 [redacted]

FIPS PUB 60-2

Supersedes FIPS PUB 60-1
1979 August 27

FEDERAL INFORMATION PROCESSING STANDARDS PUBLICATION

1983 JULY 29

U.S. DEPARTMENT OF COMMERCE/National Bureau of Standards

I/O CHANNEL INTERFACE

CATEGORY: HARDWARE STANDARD
SUBCATEGORY: INTERFACE

FIPS PUB 60-2

U.S. DEPARTMENT OF COMMERCE, Malcolm Baldrige, Secretary
NATIONAL BUREAU OF STANDARDS, Ernest Ambler, Director

Foreword

The Federal Information Processing Standards Publication Series of the National Bureau of Standards is the official publication relating to standards adopted and promulgated under the provisions of Public Law 89-306 (Brooks Act) and under Part 6 of Title 15, Code of Federal Regulations. These legislative and executive mandates have given the Secretary of Commerce important responsibilities for improving the utilization and management of computers and automatic data processing in the Federal Government. To carry out the Secretary's responsibilities, the NBS, through its Institute for Computer Sciences and Technology, provides leadership, technical guidance, and coordination of Government efforts in the development of guidelines and standards in these areas.

Comments concerning Federal Information Processing Standards Publications are welcomed and should be addressed to the Director, Institute for Computer Sciences and Technology, National Bureau of Standards, Washington, DC 20234.

James H. Burrows, Director
Institute for Computer Sciences and Technology

Abstract

This standard defines the functional, electrical, and mechanical interface specifications for connecting computer peripheral equipment as a part of automatic data processing (ADP) systems. This standard, together with a companion standard for power control, defines the hardware characteristics for the I/O channel level interface. In order to achieve full plug-to-plug interchangeability of peripheral components, device class specific operational specifications standards are also required for each class of peripheral device. These operational specifications standards will be proposed as Federal Information Processing Standards to accompany this standard as they are developed.

The Government's intent in employing this I/O Channel Interface standard is to reduce the cost of satisfying the Government's data processing requirements through increasing its available alternative sources of supply for computer system components at the time of initial system acquisition, as well as in system replacement and augmentation and in system component replacement.

This standard is also expected to lead to improved reutilization of system components. When acquiring ADP systems and system components, Federal agencies shall cite this standard in specifying the interface for connecting computer peripheral equipment as a part of ADP systems.

Key words: automatic data processing (ADP); channel level power control interface; computer peripherals; computers; Federal Information Processing Standard; input/output; interfaces.

Natl. Bur. Stand. (U.S.), Fed. Info. Process. Stand. Publ. (FIPS PUB) 60-2, 97 pages.
(1983)
CODEN:FIPPAT

For sale by the National Technical Information Service, U.S. Department of Commerce, Springfield, VA 22161.

**Federal Information
Processing Standards Publication 60-2**

1983 July 29

Announcing the Standard for

I/O CHANNEL INTERFACE

Federal Information Processing Standards Publications are issued by the National Bureau of Standards pursuant to section 111(f)(2) of the Federal Property and Administrative Services Act of 1949, as amended, Public Law 89-306 (79 Stat. 1127), Executive Order 11717 (38 FR 12315, dated May 11, 1973) and Part 6 of Title 15 Code of Federal Regulations (CFR).

Name of Standard. I/O Channel Interface (FIPS PUB 60-2).

Category of Standard. Hardware Standard, Interface.

Explanation. This standard defines the functional, electrical, and mechanical interface specifications for connecting computer peripheral equipment as a part of automatic data processing (ADP) systems. This standard, together with a companion standard for power control, defines the hardware characteristics for the I/O channel level interface. In order to achieve full plug-to-plug interchangeability of peripheral components, device class specific operational specifications standards are also required for each class of peripheral device. Three of these operational specifications standards are already approved as FIPS 62, FIPS 63-1, and FIPS 97, and others will be proposed as they are developed.

The Government's intent in employing this I/O Channel Interface standard is to reduce the cost of satisfying the Government's data processing requirements through increasing its available alternative sources of supply for computer system components at the time of initial system acquisition, as well as in system replacement and augmentation and in system component replacement.

This standard is also expected to lead to improved reutilization of system components. When acquiring ADP systems and system components, Federal agencies shall cite this standard in specifying the interface for connecting computer peripheral equipment as a part of ADP systems.

Approving Authority. Secretary of Commerce.

Maintenance Agency. Department of Commerce, National Bureau of Standards (Institute for Computer Sciences and Technology).

Cross Index. The original version of this standard was based upon American National Standards Institute document X3T9/600, Rev. 2, Draft Proposed American National Standard for I/O Channel Interface.

Applicability. This standard is applicable to the acquisition of all ADP systems and peripheral equipment for those systems except those minicomputer, microcomputer, and other small scale systems which are specifically excluded by the National Bureau of Standards (NBS). A list of such currently excluded systems and the current criteria for exclusion will be developed and maintained by NBS and will be periodically distributed to all Federal agencies and be publicly available upon request.

This standard is applicable to the acquisition of (1) all new and replacement ADP systems, (2) computer peripheral equipment acquired to replace existing peripheral equipment of or to augment ADP systems that employ interfaces conforming with this standard, and (3) peripheral equipment acquired to replace existing peripheral equipment of or to augment ADP systems that do not conform to this standard, but for which the hardware and software necessary to conform to this standard are commercially available.

Specifically, this standard shall be employed in the interconnection of computer peripheral equipment as a part of ADP systems for the following types of peripherals: (1) magnetic tape equipment employing open reel-to-reel magnetic tape storage devices, specifically excluding magnetic tape cassette and tape cartridge storage devices, (2) magnetic disk storage equipment employing disk drives each having a capacity greater than 7 megabytes per storage module, specifically excluding flexible disk and disk cartridge devices having a smaller storage capacity per device, and (3) peripheral equipment employing all peripheral device types for which accompanying operational specifications standards have been issued as Federal Information Processing Standards.

Verification of the correct operation of all interfaces that are required to conform to this standard shall, through demonstration or other means acceptable to the Government, be provided prior to the acceptance of all applicable ADP equipment.

Specifications. This standard incorporates by reference the technical specifications of the following NBS document: Technical Specifications for I/O Channel Interface, dated August 3, 1982. Copies of the technical specifications section of the standard will be available from the National Technical Information Service as described in the **Where to Obtain Copies** section below.

Implementation. The original version of this standard became effective June 23, 1980, and the provisions from it which this revision retains continue in effect from that date. The changes made by this revision become effective July 29, 1983.

All applicable equipment ordered on or after the effective date, or procurement actions for which solicitation documents have not been issued by that date, must conform to the provisions of this standard unless a waiver has been granted in accordance with the procedure described elsewhere in this standard. In addition, in the absence of such a waiver, interconnection of any type of peripheral equipment specified in the third paragraph under "Applicability," as a part of any ADP system described in the second paragraph under that heading, shall conform to this standard, FIPS 61, and the applicable peripheral standard(s) where either the peripheral equipment or the ADP system or both were acquired on or after June 23, 1980.

Regulations concerning the specific use of this standard in Federal procurement will be issued by the General Services Administration to be a part of the Federal Property Management Regulations.

This standard shall be reviewed by NBS within three years after its effective date, taking into account technological trends and other factors, to determine whether the standard should be affirmed, revised, or withdrawn.

Waivers. Heads of agencies desiring a waiver from the requirements stated in this standard, so as to acquire ADP equipment that does not conform to this standard, shall submit a request for such a waiver to the Secretary of Commerce for review and approval. Approval will be granted if, in the judgment of the Secretary based on all available information, including that provided in the waiver request, a major adverse economic or operational impact would occur through conformance with this standard.

A request for waiver shall include: (1) a description of the existing or planned ADP system for which the waiver is being requested, (2) a description of the system configuration, identifying those items for which the waiver is being requested, and including a description of planned expansion of the system configuration at any time during its life cycle, and (3) a justification for the waiver, including a description and discussion of the major adverse economic or operational impact that would result through conformance to this standard as compared to the alternative for which the waiver is requested.

The request for waiver shall be submitted to the Secretary of Commerce, Washington, D. C. 20230, and labeled as a Request for Waiver to a Federal Information Processing Standard. Waiver requests will normally be processed within 45 days of receipt by the Secretary. No action shall be taken to issue solicitation documents or to order equipment to which this standard is applicable and which does not conform to this standard prior to receipt of a waiver approval response from the Secretary.

Where to Obtain Copies. Either paper or microfiche copies of this Federal Information Processing Standard, including the technical specifications, may be purchased from the National Technical Information Service (NTIS) by ordering Federal Information Processing Standard Publication 60-2 (FIPS PUB 60-2), I/O Channel Interface. Ordering information, including prices and delivery alternatives, may be obtained by contacting the National Technical Information Service (NTIS), U. S. Department of Commerce, Springfield, Virginia 22161, telephone: (703) 487-4650.

TECHNICAL SPECIFICATIONS

FOR

I/O CHANNEL INTERFACE

FOREWORD

This standard is a revision of FIPS PUB 60-1 which incorporated a draft proposed American National Standard Specifications for I/O Channel Interface, X3T9/600 Rev. 2, dated August 18, 1976, which in turn was largely derived from IBM System/360 and System/370 I/O Interface to Control Unit Original Equipment Manufacturers' Information, IBM publication number GA22-6974-2. This revision incorporates material from a later revision of that same IBM manual, IBM publication number GA22-6974-5.

This revision contains a number of detailed changes to the text of FIPS PUB 60-1; however, most are editorial revisions for the sake of clarity. Significant technical changes include:

a. The use of the 'clock out' signal is now optional.

b. The data-streaming feature is added. This feature changes the normal dc-interlocked handshaking of data transfers to a noninterlocked mode, permitting high data transfer rates through much longer cables than would be possible in the dc-interlocked mode.

c. The sense ID command, which is recognized by recent compatible control units, has been added to the list of "generic" commands specified in the I/O channel standard.

TABLE OF CONTENTS

CHAPTER 1.

Introduction

1.1 Scope. This Federal Information Processing Standard specifies functional, electrical, and mechanical characteristics of the interface between I/O control units and channels in general purpose computer systems.

1.2 Definitions.

Channel Program - A sequence of CCW's which control the operation of a subchannel.

Channel Command Word (CCW) - A CCW specifies the information necessary for the channel, control unit, and addressed device to initiate an I/O operation. Information in the CCW includes the operation to be performed, the address and length of the relevant data, and a set of flags indicating what is to be done at the termination of the operation.

Subchannel - The channel facilities required for sustaining a single I/O operation. These facilities include registers which contain address, command, count, and status information relating to the I/O operation currently in progress.

CHAPTER 2.

Functional Specifications

Sections 2.1 through 2.8 define the functional requirements for the basic interface. Features, which can be added to the basic interface, are defined in section 2.9.

2.1 General. Input/output devices provide external storage and a means of communication between data processing systems or between a system and the external world.

Input/output (I/O) devices are attached to the central processing unit (CPU) by means of channels and control units in various configurations (figure 1).

The channel directs the flow of information between I/O control units and main storage.

The control unit provides the logical capability necessary to operate and control an I/O device and adapts the characteristics of each I/O device to the standard form of control provided by the channel. A control unit may be housed separately or it may be physically and logically integrated with the I/O device.

An I/O device attached to the control unit may be designated to perform only certain limited operations; for example, recording data and moving the recording

medium. To accomplish these functions, the I/O device needs detailed signal sequences peculiar to that type of I/O device. The control unit decodes the commands received from the channel, interprets them for the I/O device, and provides the signal sequence for executing the operation.

The connection between the channel and the control unit is called the I/O interface. The I/O interface provides an information format and a signal sequence common to all control units. The interface consists of signal lines that connect a number of control units to a channel (figure 1). Except for signals used to establish selection control, all communications to and from the channel occur over a common bus; i.e., any signal provided by the channel is available to all control units (figure 2). At any one time, however, only one control unit can be logically connected to the channel. Selection of a control unit for communication with the channel is controlled by a signal (passing serially through all control units) that permits, sequentially, each control unit to respond to the signals provided by the channel. A control unit remains logically connected to the interface until it transfers the information it needs or has, or until the channel signals it to disconnect.

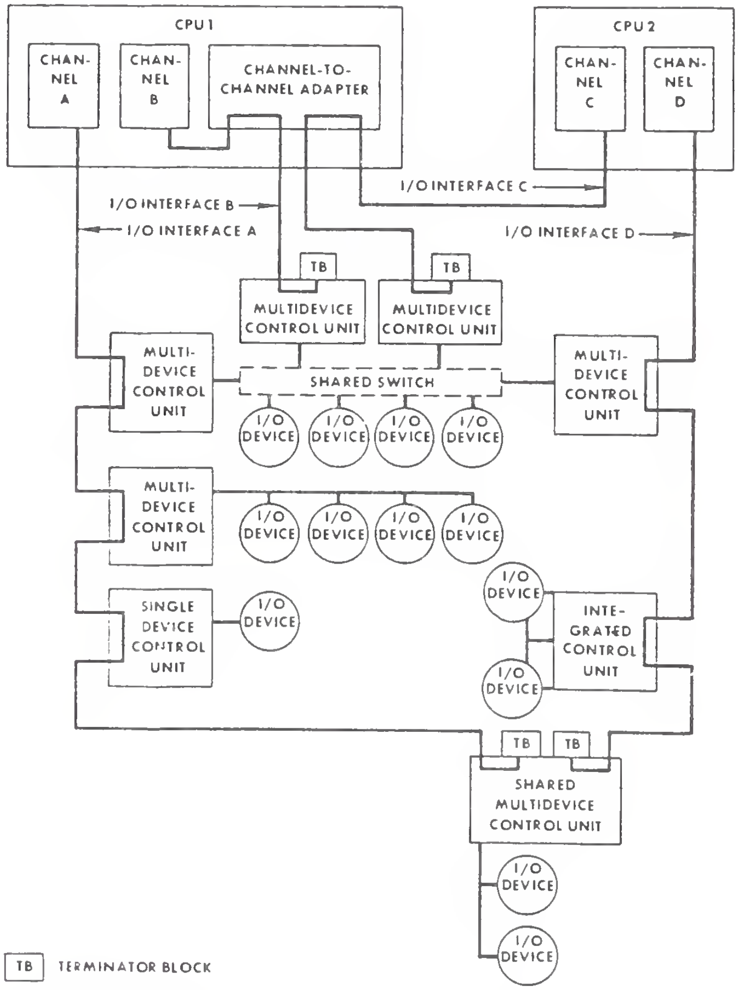

FIGURE 1. Input/Output Interface—Multiple Configurations

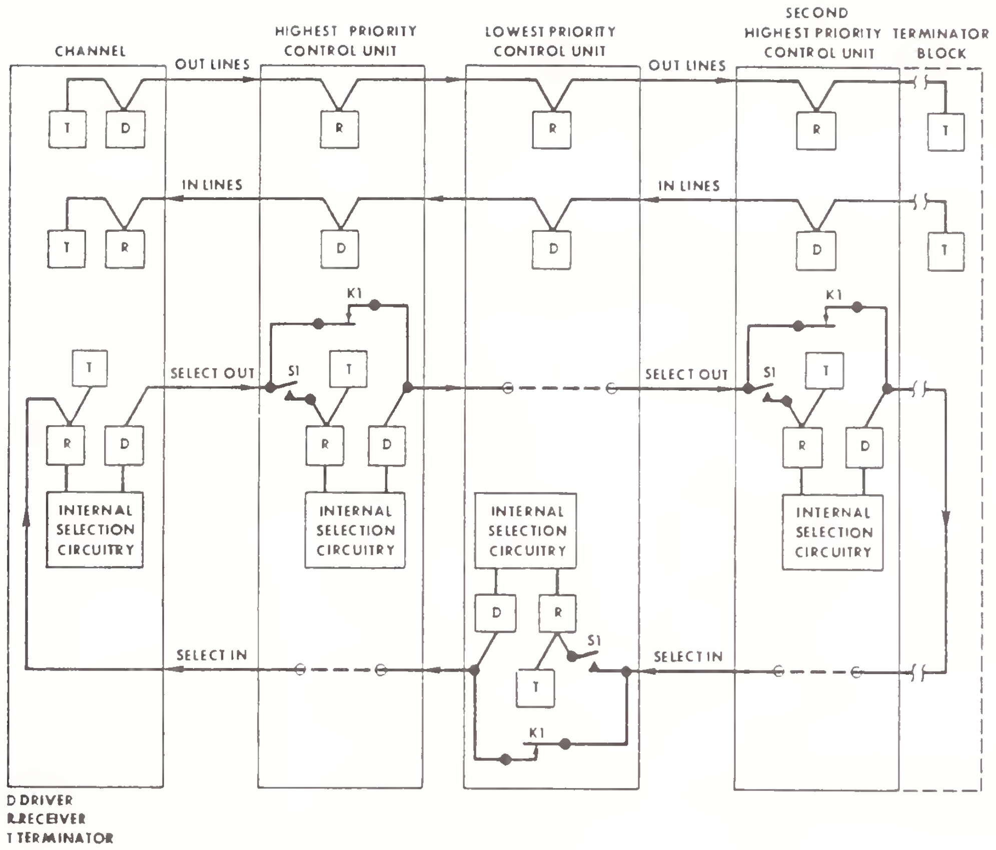

FIGURE 2. Interconnections on the Input/Output Interface

The interface can accommodate up to 256 directly addressable I/O devices (limit set by addressing facilities). The number of control units that can be accommodated is limited only by timing and electrical considerations within the given addressing limitations. (See section 3.1 "Multiple Drivers and Receivers" and section 2.8.3 "System Configuration.")

The multiplexing facilities of the interface permit the possibility of operating any number of the 256 I/O devices concurrently on a single interface; portions of various messages can be transmitted over the interface in an interleaved fashion to or from different I/O devices, or a complete message can be transmitted in a single interface operation. The operation is determined by the particular channel and control unit.

The rise and fall of all signals transmitted over the interface are generally controlled by dc-interlocked responses. The dc-interlocked sequences remove the dependence of the interface on circuit speed, and make it applicable to a wide variety of circuits and data rates. Further, dc-interlocking permits connecting control units of different circuit speeds to a single channel. Those sequences which are not dc-interlocked allow for designs capable of achieving higher data rates than

those achievable through dc-interlocked sequences. However, the electrical specifications of circuits not using dc-interlocked sequences are more restrictive and do not allow as wide a range of possible circuit speeds.

2.2 Line Definition. The I/O interface connects a channel with control units. External cables physically connect all control units in a chain, with the first control unit connected to the channel. (See figures 1 and 2.)

Input/Output Interface Lines

The I/O interface lines and their uses are:

Name of Line	Abbreviation	Uses
Bus 0 Out Position P	Bus 0 Out P	BUS OUT—Used to transmit information (data, I/O device address, command, control orders) from the channel to the control unit. ('Bus 1' is only available with the bus extension feature.)
Bus 0 Out Position 0	Bus 0 Out 0	
Bus 0 Out Position 1	Bus 0 Out 1	
Bus 0 Out Position 2	Bus 0 Out 2	
Bus 0 Out Position 3	Bus 0 Out 3	
Bus 0 Out Position 4	Bus 0 Out 4	
Bus 0 Out Position 5	Bus 0 Out 5	
Bus 0 Out Position 6	Bus 0 Out 6	
Bus 0 Out Position 7	Bus 0 Out 7	
Bus 1 Out Position P	Bus 1 Out P	
Bus 1 Out Position 0	Bus 1 Out 0	
Bus 1 Out Position 1	Bus 1 Out 1	
Bus 1 Out Position 2	Bus 1 Out 2	
Bus 1 Out Position 3	Bus 1 Out 3	
Bus 1 Out Position 4	Bus 1 Out 4	
Bus 1 Out Position 5	Bus 1 Out 5	
Bus 1 Out Position 6	Bus 1 Out 6	
Bus 1 Out Position 7	Bus 1 Out 7	
Bus 0 In Position P	Bus 0 In P	BUS IN—Used to transmit information (data, selected I/O device identification, status information, sense data) from the control unit to the channel. ('Bus 1' is only available with the bus extension feature.)
Bus 0 In Position 0	Bus 0 In 0	
Bus 0 In Position 1	Bus 0 In 1	
Bus 0 In Position 2	Bus 0 In 2	
Bus 0 In Position 3	Bus 0 In 3	
Bus 0 In Position 4	Bus 0 In 4	
Bus 0 In Position 5	Bus 0 In 5	
Bus 0 In Position 6	Bus 0 In 6	
Bus 0 In Position 7	Bus 0 In 7	
Bus 1 In Position P	Bus 1 In P	
Bus 1 In Position 0	Bus 1 In 0	
Bus 1 In Position 1	Bus 1 In 1	
Bus 1 In Position 2	Bus 1 In 2	
Bus 1 In Position 3	Bus 1 In 3	
Bus 1 In Position 4	Bus 1 In 4	
Bus 1 In Position 5	Bus 1 In 5	

Bus 1 In Position 6	Bus 1 In 6	
Bus 1 In Position 7	Bus 1 In 7	
Mark 0 In	Mk 0 In	MARK—Used to indicate the bus(es) being used. (Except for 'Mark 0 In,' these lines are only used with the bus extension feature.)
Mark 0 Out	Mk 0 Out	
Mark 1 In	Mk 1 In	
Mark 1 Out	Mk 1 Out	
Mark In Parity	Mk In P	
Mark Out Parity	Mk Out P	
Address Out	Adr Out	TAGS—Used for interlocking and controlling information on the buses and for special sequences. ('Disconnect In' is only used with the I/O error alert feature. 'Data In' and 'Data Out' are only used with the high-speed transfer and data streaming features.)
Address In	Adr In	
Command Out	Cmd Out	
Status In	Sta In	
Service Out	Srv Out	
Service In	Srv In	
Data In	Dat In	
Data Out	Dat Out	
Disconnect In	Dis In	
Operational Out	Opl Out	SELECTION CONTROLS—Used for the scanning of, or the selection of, attached I/O devices.
Operational In	Opl In	
Hold Out	Hld Out	
Select Out	Sel Out	
Select In	Sel In	
Suppress Out	Sup Out	
Request In	Req In	
Metering Out	Mtr Out	METERING CONTROLS—Used for the conditioning of usage meters located in the various attached units.
Metering In	Mtr In	
Clock Out	Clk Out	

Note: Except when the data streaming feature is used, the validity of information on the buses and the timing of the signals on the tag lines is specified at the channel cable connectors.

2.2.1 Buses. Each bus is a set of nine lines consisting of eight information lines and one parity line. Information on the buses is arranged so that bit position 7 of a bus always carries the low-order bit within an eight-bit byte. The highest-order bit is in position 0 and intervening bits are in descending order from position 1 to position 6.

When a byte transmitted over the interface consists of less than eight information bits, the bits must be placed in the highest-numbered contiguous bit positions of the bus. Any unused lines of the bus must include the low-numbered positions (position 0 and adjacent positions). Unused lines must present logical zeros to the receiving end. The parity bit of any byte must appear in the parity position (P). The byte must always have odd parity (figure 3).

Bus Line	BCD (Position Value)	Packed Numeric (Position Value)		Unpacked Numeric (Position Value)	EBCDIC (Bit Positions)	ASCII (Bit Positions)	Binary (Position Value)
P	P	P		P	P	P	P
0	0	8		0	0	(zero vaue)	128
1	0	4	Digit x	0	1	7	64
2	B	2		0	2	6	32
3	A	1		0	3	5	16
4	8	8		8	4	4	8
5	4	4	Digit x + 1	4	5	3	4
6	2	2		2	6	2	2
7	1	1		1	7	1	1

X = Higher order digit
X + 1 = Lower order digit
0 = Logical Zero

FIGURE 3. Organization of Information on Bus

2.2.1.1 Bus Out. 'Bus out' is used to transmit addresses, commands, control orders, and data to the control units. The type of information transmitted over 'bus out' is indicated by the outbound tag lines:

1. Except for interface disconnect, when 'address out' is up during the channel-initiated selection sequence, 'bus out' specifies the address of the I/O device with which the channel wants to communicate (see section 2.5.10 "Interface Disconnect").

2. When 'command out' is up in response to 'address in' during the channel-initiated selection sequence, 'bus out' specifies a command.

3. When 'service out' is up in response to 'service in' (or 'data out' is up in response to 'data in') during the execution of a write or control command, 'bus out' specifies a data byte that is being transferred between the channel and the control unit. For example, during an operation that is specified by a write command, 'bus out' contains data to be recorded by the I/O device. During an operation that is specified by a control command, 'bus out' may contain either data that describes in more detail how the operation is to be executed or a second-level address within the control unit or I/O device.

Note: When the data-streaming feature is used, the following paragraph does not apply. (See section 2.9.5 "Data-Streaming Feature".)

When dc-interlocking is used, the period during which information on 'bus out' is valid is controlled by the tag lines. During transmission of the I/O-device address during the initial-selection sequence, information on 'bus out' must be valid from the rise of 'address out' until the rise of 'operational in' or 'select in.' In the short-busy sequence, the I/O-device address on the bus-out line must be valid from the rise of 'address out' until the fall of 'status in.' When the channel is transmitting any other type of information, the information on 'bus out' is valid from the rise of the signal on the associated outbound tag line until the fall of the signal on the responding inbound tag line.

Some skew on 'bus out' is accommodated by the channel. The channel provides a delay that accommodates skew caused by its own circuitry and, in addition, provides a delay of at least 100 nanoseconds. This delay compensates for skew caused by the cable and, for most control units, also is sufficient to accommodate the skew caused by the I/O-interface receivers. (Except as noted under "Address Out" in this chapter, the channel delays raising of the signal on the outbound tag lines so that the information on 'bus out' precedes the signal on the outbound tag line by at least 100 nanoseconds. This delay is measured at the channel cable connectors and under the worst-case skew conditions.) The control unit, when it can cause more skew, provides additional delay to compensate for the skew.

2.2.1.2 Bus In. 'Bus in' is used to transmit addresses, status, and data to the channel. The type of information transmitted over 'bus in' is indicated by the inbound tags. A control unit can place and maintain information on 'bus in' only when its 'operational in' is up, except in the short-busy sequence. (See section 2.3.2 "Short-Busy Sequence.")

The type of information transmitted over 'bus in' is indicated by the inbound tag lines:

1. When 'address in' is up, 'bus in' specifies the address of the currently selected I/O device.

2. When 'status in' is up, 'bus in' contains a byte of information that describes the status of the I/O device or control unit.

3. When 'service in' (or 'data in') is up during execution of a read, read-backward, or sense command, the nature of the information contained on 'bus in' depends on the type of operation. During a read operation, it may contain a byte of data from the record medium. During a sense operation, 'bus in' contains data describing unusual conditions detected at the I/O device.

Note: When the data-streaming feature is used, the following paragraph does not apply. (See section 2.9.5 "Data-Streaming Feature.")

The period during which information on 'bus in' is valid is controlled by the tag lines. Information on the bus is valid within 100 nanoseconds after the rise of the

associated inbound tag and remains valid until the rise of the responding outbound tag or, in a short-busy sequence, until the fall of 'select out.' The 100-nanosecond delay between the rise of the inbound tag and the time the signal becomes valid on 'bus in' places the responsibility on the channel for deskewing 'bus in.' The channel provides a delay in the inbound tag lines to accommodate skew caused by the channel circuitry (including its receivers) and, in addition, provides a delay of at least 100 nanoseconds. This delay compensates for skew caused by the cable, and, for most control units, for the skew caused by their drivers. This delay provides sufficient time to deskew the information so that the inbound tag can be raised by the control unit at the same time information is placed on the bus. When a control unit and cable can cause more skew than can be accommodated by a 100-nanosecond delay, the control unit provides the additional delay to eliminate this greater skew.

2.2.2 Selection Controls and Tag Lines.

2.2.2.1 Operational Out. 'Operational out' is a line from the channel to all attached control units and is used for interlocking purposes. Except for 'suppress out,' all lines from the channel are significant only when 'operational out' is up. Whenever 'operational out' is down, all inbound lines from the control unit drop, and any operation currently in process over the interface is reset. Under these conditions, all control-unit-generated interface signals are down within 1.5 microseconds after the fall of 'operational out' at the control unit. (See section 2.5.11 "Selective Reset" and section 2.5.12 "System Reset.")

2.2.2.2 Request In. 'Request in' is a line from all attached control units to the channel. This line, when raised, indicates that the control unit requires service and is requesting a selection sequence.

'Request in' is dropped when:

1. 'Operational in' rises, unless additional control-unit-initiated sequences are required, or

2. The control unit is no longer ready to present the status information or data, or

3. The selection requirement is satisfied by another path.

'Request in' never falls later than 250 nanoseconds after the fall of 'operational in' if the sequence satisfies the service requirements of the control unit.

'Request in' does not remain up when 'suppress out' is up if the request is for presentation of suppressible status. (See section 2.5.7 "Suppress Status.") When the control unit is requesting a selection sequence in order to present suppressible status, 'request in' falls at the control unit within 1.5 microseconds after the rise of 'suppress out' at the control unit.

'Request in' can be signalled by more than one control unit at a time.

2.2.2.3 Address Out. 'Address out' is a tag line from the channel to all attached control units. 'Address out' is used to signal all the control units to decode the

I/O-device address on 'bus out.' If the control unit recognizes the address, it responds by raising 'operational in' when 'select out' ('hold out') rises with 'address out' still up (except in the short-busy sequence). (See section 2.2.2.4 "Select Out/Hold Out and Select In.") 'Address out' rises 250 nanoseconds after the I/O-device address is placed on 'bus out' or after the rise of 'operational out,' whichever occurs later. 'Address out' is down for at least 250 nanoseconds before its rise for I/O-device selection. If 'address out' falls before 'select out' rises, the I/O-device selection is canceled.

'Address out' can rise only when 'select out' ('hold out'), 'select in,' 'status in,' and 'operational in' are down at the channel. Ultimate use of the I/O-device address on 'bus out' at the control unit is timed by the next rise of 'select out' at the addressed control unit. The rise of 'address out' is delayed at least 250 nanoseconds after the address is placed on 'bus out.' Once 'address out' and 'select out' ('hold out') are up, 'address out' stays up until either 'select in' or 'operational in' rises or, in the short-busy sequence, until 'status in' falls. Except when interface disconnect is being signalled, during I/O-device selection, 'address out' cannot be up concurrently with any other outbound tag line. (See section 2.5.10 "Interface Disconnect.")

2.2.2.4 Select Out/Hold Out and Select In. Control-unit selection is controlled by 'select out,' select in,' and 'hold out.' 'Select out' and 'select in' form a loop from the channel through each control unit to the cable terminator block ('select out') and again through each control unit back to the channel ('select in'). Control-unit-selection circuitry may be attached to either 'select out' or 'select in.' In this manual, the selection circuitry of all control units is assumed to be attached to 'select out.' All discussions that apply to the selection logic of control units attached to 'select out' equally apply in cases where the control unit is attached to 'select in.' The selection priority is (1) all control units with selection circuitry attached to 'select out,' in order of attachment from the channel to the cable terminator block, followed by (2) all control units with selection circuitry attached to 'select in,' in order of attachment from the cable terminator block to the channel. (See figure 2.) If selection is not required, the selection signal is in turn propagated by each control unit to the next control unit on the line.

Each control unit ensures that the process of electrically bypassing 'select out' before power changes does not interfere with the propagation of 'select out.' Thus, 'select out' discontinuities, which may occur when another control unit on the I/O interface is powered up or down, do not affect the propagation of 'select out.' This protection should be accomplished by the use of a special latch circuit. The latch is turned on by the AND of 'select out' and 'hold out' and is reset by the fall of 'hold out.' The circuit is in series with the remaining selection circuitry in the control unit and provides a constant 'select out' within the control unit--and therefore to the following control unit--regardless of variations in the input 'select out' signal. (See figure 4.)

Throughout the following description, 'select out' assumes the latch operates properly, that is, the rise of 'select out' at the control unit assumes that 'hold out' is up and that the fall of 'select out' is a result of the fall of 'hold out.'

A control unit can raise its 'operational in' only at the rise of its incoming 'select out.' Once a control unit propagates 'select out,' it cannot raise its 'operational in'

or respond with a short-busy sequence until the next rise of the incoming 'select out.'

When an operation is being initiated by the channel, 'select out' is raised not less than 400 nanoseconds after the rise of 'address out,' which indicates the address of the device being selected.

The channel keeps 'select out' up until 'select in' rises, or 'address in' and 'operational in' rise, or 'status in' rises.

When 'select in' rises, 'select out' drops and does not rise again until after 'select in' falls.

A control unit becomes selected only when it raises its 'operational in.' After the drop of 'select out,' the control unit keeps 'operational in' up until the current signal sequence is completed. If a control unit raises 'operational in,' it suppresses the propagation of 'select out' to the next control unit. If the control unit does not require selection, it propagates 'select out' to the next control unit within 1.8 microseconds. (See section 2.8.1 "Interface Timeout Considerations.")

When 'status in' rises in response to 'select out' in the short-busy sequence, 'select out' drops and does not rise until 'address out' has dropped.

'Hold out' is a line from the channel to all attached control units and is used in conjunction with 'select out' to synchronize control-unit selection.

'Hold out' is also used to minimize the propagation of the fall of 'select out' by purging the 'select out' signal from the 'select out' signal path. Therefore, once 'hold out' drops, it does not rise for at least 4 microseconds in general system configurations. The minimum downtime of this signal may be optionally adjusted at installation time to a minimum of 2 microseconds to handle high-speed channel configurations. In all cases, the channel is capable of providing the 4-microsecond timing for general system configurations. (See section 2.8.3 "System Configuration.")

To prevent overlapping of interface sequences, one of the following procedures is performed:

1. 'Select out' is not raised until all inbound signals for the preceding sequence are in a down state.

2. In-tags are not considered valid until 1.5 microseconds after the fall of 'operational in' for the preceding sequence.

'Select in' is a line that extends the select out signal from the jumper in the cable terminator block to the channel (see figure 2). It provides a return path to the channel for the 'select out' signal.

2.2.2.5 Operational In. 'Operational in' is a line from all attached control units to the channel and is used to signal the channel that an I/O device has been selected. 'Operational in' stays up for the duration of the selection. The selected I/O device is identified by the address byte transmitted over 'bus in' when 'address in' was raised.

The rise of 'operational in' indicates that an I/O device is selected and the control unit is actively communicating with the channel. 'Operational in' rises oly when the incoming 'select out' to the control unit is up and the outgoing 'select out' is down; that is, the control unit raises 'operational in' (except in the short-busy sequence) only in response to the rise of 'select out' and blocks the 'select out' signal from being propagated to the next control unit. 'Operational in' drops only after 'select out' drops.

When 'operational in' is raised for a particular signal sequence, it stays up until all required information is transmitted between the channel and the control unit. If 'select out' is down, 'operational in' drops after the rise of the outbound tag associated with the transfer of the last byte of information. With the exception of 'request in' or 'metering in,' all inbound signals are down within 100 nanoseconds after the fall of 'operational in' at the control unit.[1]

2.2.2.6 Address In. 'Address in' is a tag line from all attached control units to the channel and is used to signal the channel when the address of the currently selected I/O device has been placed on 'bus in.' During an initial-selection sequence or a control-unit-initiated sequence, the channel responds to 'address in' by raising 'command out.' 'Address in' stays up until the rise of 'command out.' 'Address in' must fall so that 'command out' may fall. 'Address in' is not up concurrently with any other inbound tag line.

2.2.2.7 Command Out. 'Command out' is a tag line from the channel to all attached control units and is used to signal the selected I/O device in response to a signal on 'address in,' 'status in,' 'data in,' or 'service in.' The rise of 'command out' indicates that any information on 'bus in' is no longer required to be valid. 'Command out' stays up until the fall of the associated 'address in,' 'status in,' 'data in,' or 'service in.' However, when the data-streaming feature is used, 'command out' is not dc-interlocked with 'service in' or 'data in.' (See section 2.9.5 "Data-Streaming Feature.") 'Command out' is not up concurrently with any other outbound tag line, except during an interface-disconnect sequence, when 'address out' may be up. (See section 2.5.10 "Interface Disconnect.")

During an initial-selection sequence, 'command out' rising in response to the rise of 'address in' indicates to the selected I/O device that the channel has placed a command byte on 'bus out.' (The command byte has a fixed format. See section 2.4 "Commands.") 'Command out' in response to 'data in' or 'service in' always means stop. (See section 2.5.2 "Stop".) 'Command out' in response to 'status in' means stack. (See section 2.5.3 "Stack Status.")

---

[1]For some older control units and devices designed, all inbound signals except 'request in' and 'metering in' are down within 1.5 microseconds after the fall of 'operational in' at the control unit.

When 'command out' is raised to indicate proceed, stack, or stop, 'bus out' has a byte of all zeros but does not necessarily have correct parity. 'Bus out' is not checked for parity or decoded as a command by the control unit under these circumstances.

2.2.2.8 Status In. 'Status in' is a tag line from all attached control units to the channel and is used to signal the channel when the selected control unit has placed status information on 'bus in.' The status byte has a fixed format and contains bits describing the current status at the control unit. (See section 2.6.1 "Status Byte.") The channel responds by raising either 'service out' or 'command out' or, in the case of the short-busy sequence, by dropping 'select out.'

'Status in' is not up concurrently with any other inbound tag line. 'Status in' stays up until the rise of an out tag or, in the short-busy sequence, until 'select out' falls. 'Status in' must fall so that the responding out tag may fall. In the short-busy sequence, status information on 'bus in' is valid until 'select out' ('hold out') falls.

2.2.2.9 Service Out. 'Service out' is a tag line from the channel to all attached control units and is raised to signal the selected I/O device when 'service in' or 'status in' has been recognized. A signal on 'service out' indicates to the selected I/O device that the channel has accepted the information on 'bus in' or has provided on 'bus out' the data requested by 'service in.'

Note: When the data-streaming feature is used, the following two paragraphs do not apply. (See section 2.9.5 "Data-Streaming Feature.")

When 'service out' rises in response either to 'service in' during a read, read-backward, or sense operation or to 'status in,' the information placed on 'bus in' by the control unit has been accepted by the channel. In these cases, the rise of 'service out' indicates that the information on 'bus in' is no longer required to be valid and is not associated with any informaton on 'bus out.' When 'service out' rises in response to 'service in' during a write or control operation, the rise of 'service out' indicates that the channel has placed the requested information on 'bus out.'

'Service out' stays up until the fall of the associated 'service in' or 'status in.' 'Service out' is not up concurrently with any other out tag except during an interface-disconnect sequence, when 'address out' may be up. (With the high-speed transfer feature, 'data out' and 'service out' may be up concurrently.)

A 'service out' response to 'status in' while 'suppress out' is up indicates to the control unit that the status is accepted and the operation is to be chained to a new operation. (See section 2.5.9 "Command Chaining.")

2.2.2.10 Service In. 'Service in' is a tag line from all attached control units to the channel and is used to signal to the channel when the selected I/O device is ready to send or receive a byte of information. The nature of the information associated with 'service in' depends on the operation and the I/O device. The channel responds to the rise of 'service in' by raising either 'service out' or 'command out.'

Note: When the data-streaming feature is used, the following two paragraphs do not apply. (See section 2.9.5 "Data-Streaming Feature.")

During read, read-backward, and sense operations, 'service in' rises when information is available on 'bus in.' During execution of operations specified by either a write or control command, 'service in' rises when information is required on 'bus out.' 'Service in' is not up concurrently with any other inbound tag line. 'Service in' stays up until the rise of either 'service out,' 'command out,' or 'address out.' (With the high-speed transfer feature, 'data in' and 'service in' may be up concurrently.)

When, in the case of cyclic I/O devices, the channel does not respond in time to the preceding 'service in,' and overrun condition may be recognized by the control unit or I/O device. In this situation, data transfer is terminated, and the unit-check status indicator and the overrun sense indicator are set to ones. When this condition is recognized, 'service in' does not drop if an out-tag has not risen and does not rise if 'service out' has not dropped.

For I/O devices that may overrun, the critical signal timings involved are given in the vendor's documentation for the I/O device.

2.2.2.11 Suppress Out. 'Suppress out' is a line from the channel to all attached control units; it may rise or fall at any time. This line is used both alone and in conjunction with the out-tag lines to provide the following special functions: suppress data, suppress status, command chaining, and selective reset. Each of these functions is described in subsequent sections.

2.2.3 Metering Controls.

2.2.3.1 Clock Out. This line not used; the channel ensures that the 'clock out' line remains down at all times.[2]

2.2.3.2 Metering In. 'Metering in' is a line from all attached control units to the channel. The interpretation of the 'metering in' signal as provided by the control unit is defined in the vendor's manual for the system. The 'metering in' signal originates from each I/O device and/or control unit and is generated by the I/O device from the time of acceptance of a command until the generation of device end for that command. 'Metering in' possibly is not activated between the acceptance of channel end and the generation of the associated device end if the intervening time is less than 400 milliseconds. 'Metering in' also is raised concurrently with 'operational in' for any I/O interface signalling sequence that does not involve an operation (such as test I/O or a control-unit-initiated sequence for synchronous status presentation). If 'metering in' is raised, the duration of the signal does not exceed that of 'operational in.' 'Metering in' may be signalled by more than one control unit at a time. Refer to the vendor's manual pertinent to the I/O device.

---

[2]For many older control units, the 'clock out' line from the channel to these control units is used to provide the CPU interlock control for changing the enable/disable states of the units (the signal must be down to permit changing states). In addition, the control-unit transition between the enabled and disabled state requires the same conditions as prevail for the offline/online transition. The down state of 'clock out' is at least 1 microsecond in duration. (See section 2.8.4 "Offline/Online.")

'Metering in' is not raised:

1. Between the generation and acceptance of device end.

2. Between the generation of device end and the acceptance of the next command during command chaining.

3. While a device is waiting for initiation of an automatic start. For example, transmission control units do not necessarily activate 'metering in' during the idle portion of prepare commands.

2.2.3.3 Metering Out. 'Metering out' is a line from the channel to all attached control units and is used to condition meters when they are present in assignable control units and I/O devices. The vendor's manual for the system defines when 'metering out' is activated.

2.2.4 Reserved Lines. Some of the signal lines in the I/O interface are reserved. (See section 3.5.5 "Interface-Connector Pin Assignments.")

All control units carry through the 48 assigned signal transmission lines in cables 1 and 2, with the exception of 'select out' (or 'select in' if low priority), from the IN cable connector to the OUT cable connector.[3] All control units using 'bus 1' carry through all 72 signal transmission lines in cables 1, 2, and 3. (See section 3.5.5 "Interface-Connector Pin Assignments" and figure 12.)

Note: Lines not terminated in any string of interface cables may cause interference from electrostatic discharge (ESD). To prevent such interferences, there should be electrical continuity in all cable conductors in each string of cables (from channel to its terminator). This electrical continuity can be achieved by using either 20-conductor cables or 24-conductor cables in a string, but not both types in the same string.

2.2.5 Signal-Interlock Summary. The following rules for direct-current interlocking of signals shall be followed in the design of channels and control units using this I/O interface. These rules describe the protocols to be followed by channels and control units in the absence of error conditions. That is, in situations where timing restrictions are specified in this manual and either the channel or control unit has failed to respond by dropping or raising the appropriate signal line within the required time interval, then it is assumed that the interlocking rules do not apply.

Note: When either the high-speed-transfer feature or the data-streaming feature is used, rules 1, 2, and 3 do not apply.

1. Except when interface disconnect is signalled, no more than one out-tag is up at any given time. During interface disconnect, 'address out' may be up with another out-tag.

---

[3]Some older control units and devices carry through the 40 assigned signal transmission lines in cables 1 and 2, except for 'select out' (or 'select in' if low priority), from the IN cable connector to the OUT cable connector.

2. No more than one in-tag is up at any given time.

3. An in-tag rises only when all tags are down, except in short-busy sequence.

Note: When the data-streaming feature is used, rules 4 and 5 do not apply.

4. An in-tag falls only after the rise of a responding out-tag, except for 'status in' in the short-busy sequence.

5. 'Service out,' 'command out,' and 'data out' rise only in response to the rise of an in-tag.

6. 'Address out' for an initial-selection sequence rises when 'select in' and 'select out' are down at the channel.

7. Once 'address out' and 'select out' have risen for an initial-selection sequence, 'address out' stays up until after the rise of 'select in' or 'operational in' or the fall of 'status in.'

8. Once 'address out' has risen for the interface-disconnect sequence, it does not drop until 'operation in' drops.

9. None of the out lines, except 'suppress out,' have meaning when 'operatonal out' is down.

10. 'Select out' rises only if 'operational in' and 'select in' are down.

11. 'Operational in' does not fall until one of the following events occurs:

    a. 'Select out' falls, and an out-tag response is sent for the last in tag of any given signal sequence.

    b. 'Operational out' falls.

    c. An interface-disconnect sequence is given.

12. 'Operational in' does not rise unless 'operational out' is up and, if currently up, drops when 'operational out' drops.

Note: Designers should carefully consider the effects of interface signal-transition times. Although transition time should not generally be a problem, some cases may exist, because of wide variations in circuit tolerance or in the physical integration of channel and control units, in which the transition time must be considered. Because signalling on the I/O interface usually causes the receiving unit to signal a response, the general rule is that a unit should not signal a response until it has fully recognized internally the receipt of any previously transmitted signal.

2.3 Operation. This section describes the detailed signal sequences for a complete I/O operation. These include: selection, data transfer, and ending sequences. For further reference, see Flow Diagrams 1-8 in Appendix A and Sequence Charts 1-5 in Appendix B.

2.3.1 Initial-Selection Sequence. To initiate an I/O operation, the channel places the address of the desired I/O device on 'bus out' and raises 'address out.' Each control unit connected to the channel attempts to decode the address on 'bus out'; however, only one control unit should be assigned to a given I/O address. To be acceptable, the address must have correct parity.

The channel then raises 'select out,' and the incoming 'select out' signal appears at the control unit for the addressed I/O device. If the control unit and the addressed I/O device are available to execute the operation, the control unit blocks the propagation of the 'select out' signal and raises the 'operational in' line. When 'operational in' rises, the channel responds by dropping 'address out.' The control unit places the address of the I/O device on 'bus in' and, after 'address out' falls, raises 'address in.' 'Hold out' with 'select out' may drop any time after 'address in' rises. After the channel checks the I/O-device address on 'bus in,' it responds by placing the command on 'bus out' and raises 'command out.' The selected control unit processes the command and drops 'address in,' which allows 'command out' to fall. After 'command out' falls, the control unit places status information on 'bus in' and raises 'status in'; this is referred to as initial status. At this time, the status is analyzed to determine if the command has been accepted. The command is considered to have been accepted if: (1) command retry is requested or (2) the command sent to the addressed I/O-device is not test I/O, and the status is zero or (3) the command sent to the addressed I/O device is not test I/O, and the status contains channel end but does not contain attention, control-unit end, busy, unit check, or unit exception. Once the command is accepted, the I/O operation is considered to have been initiated.

If the channel accepts the initial status, it responds by raising 'service out,' allowing the control unit to drop 'status in.' If the channel does not accept the initial status, it responds by raising 'command out,' allowing the control unit to drop 'status in.' When 'status in' falls, in either case, the initial-selection sequence is completed. (See "Stack Status" in this chapter.)

If, during the initial-selection sequence, the I/O device or the path to the I/O device is not available to execute the operation because it is currently being used to execute a previously initiated operation, or if the control unit is not available because it (1) is executing a previously initiated operation, or (2) has pending status for another I/O device, the control unit may signal that a busy condition exists. The control unit can indicate the busy condition in either of two ways, depending on the design of the control unit: it may continue execution of the initial-selection sequence until completion, or it may modify the sequence and cause execution of the short-busy sequence. If the control unit executes the initial-selection sequence, the busy condition is presented as initial status. If instead it executes the short-busy sequence, the busy condition is also presented. (See "Short-Busy Sequence," described next, and "Busy" in this chapter.)

Note: A channel response of 'command out' to 'status in' cannot prevent the execution of an immediate command. (See section 2.4.1.1 "Immediate Operation.")

2.3.2 Short-Busy Sequence. This sequence is initiated when the channel begins an initial-selection sequence and the control unit, because of a busy condition, prematurely ends the selection sequence. Specifically, when the channel has raised

'select out' after having placed the I/O-device address on 'bus out' and raised 'address out,' each control unit attempts to decode the address on 'bus out.' When 'select out' rises at the addressed control unit, the control unit blocks the propagation of 'select out,' places the busy status byte on 'bus in,' and raises 'status in.' 'Operational in' is not raised.

The channel signals that the status byte is no longer needed on 'bus in' by dropping 'select out.' The control unit responds by dropping 'status in' and disconnecting from the interface. The channel keeps 'address out' up until 'status in' drops, thus completing the short-busy sequence.

Note: The short-busy sequence is not used in response to an initial-selection sequence addressed to a device for which chaining has just been indicated.

2.3.3 Control-Unit-Initiated Sequence. When a control unit requires service, it signals the channel by raising 'request in.' The next time 'select out' rises at the control unit requiring service and no selection is being attempted by the channel ('address out' down), the control unit blocks the propagation of 'select out,' places the address of the I/O device on 'bus in,' and raises both 'address in' and 'operational in.' When the channel recognizes the address, 'command out' is sent to the control unit, indicating proceed. When 'command out' rises, the control unit responds by dropping 'address in,' thus completing the control-unit-initiated sequence.

If the service request is for data, the sequence proceeds as described in "Data-Transfer Sequence" below.

If the service request is for status information, the sequence proceeds as described in "Ending Sequence" in this chapter.

2.3.4 Data Transfer.

2.3.4.1 Data-Transfer Sequence. Data transfer may be requested by a control unit after an initial-selection sequence is executed. The direction of data transfer over the I/O interface is determined by the command that was passed to the control unit during that sequence.

Note: When the data-streaming feature is used, the following two paragraphs do not apply. (See section 2.9.5 "Data-Streaming Feature.")

To transmit data to the channel, the control unit places a data byte on 'bus in' and raises 'service in'; the tag and the validity of 'bus in' are maintained until the appropriate outbound tag is raised in response. When 'service out' rises, the control unit responds by dropping 'service in,' thus completing the data-transfer sequence. After 'service in' falls, the channel responds by dropping 'service out.'

To request data from the channel, 'service in' is raised, and the channel places the data on 'bus out' and signals the control unit by raising 'service out.' When 'service out' rises, the control unit responds by dropping 'service in,' thus completing the data-transfer sequence. The channel maintains the validity of 'bus out' until 'service in' falls. After 'service in' falls, the channel responds by dropping 'service out.'

After selection, as a result of either an initial-selection sequence or a control-unit-initiated sequence, the control unit may remain connected to the channel for the duration of the transfer of information. The transferred information can be a single byte of data, a byte of status, a new command, a string of data bytes, or, in the case of an immediate operation, the specification of a complete operation from initiation to reception of ending status.

2.3.4.2 I/O-Interface Connection. A connection begins at the time 'select out' rises at the control unit for the purpose of executing any sequence or sequences. The duration of the connection is under control of both the channel and the control unit. To provide a channel with a method of controlling the duration of the connection, a control unit does not disconnect from the I/O interface before 'select out' ('hold out') falls. However, the control unit may preserve the connection after the channel permits the control unit to disconnect—'select out' ('hold out') down--by holding up 'operational in.' In this manner, the control unit can extend the duration of the connection. The connection is considered to be ended when 'operational in' is dropped.

Depending on the duration of the connection, one of two modes is established: byte-multiplex or burst. (These modes are established to assist in the scheduling of concurrent execution of multiple I/O operations.)

2.3.4.3 Byte-Multiplex Mode. A byte-multiplex-mode connection is defined as any connection where the time contributed by the control unit during the connection (because of the control-unit circuitry or the sequencing method used) is equal to or less than 32 microseconds. The connection may occur for the purpose of initiating or continuing execution of an I/O operation or presenting status. The byte-multiplex mode is the normal mode for low-speed I/O devices; however, all I/O devices are designed to operate in burst mode when required by the channel.

2.3.4.4 Burst Mode. If the the time contributed by the control unit to the connection (because of the control-unit circuitry or the sequencing method used) is greater than 32 microseconds, the connection is considered burst mode. This mode is the normal mode of operation for high-speed I/O devices. These devices force burst mode (by holding up 'operational in') when attached to a channel capable of byte-multiplex operation.

Medium-speed or buffered I/O devices, which may normally operate in either mode, are equipped with a manual or programmable switch to select the mode of operation. The switch setting is overridden when burst mode is forced by the channel. Whenever an interface disconnect, selective reset, or system reset is executed, the force-burst-mode condition of a control unit is overridden. (See "Interface Disconnect," "Selective Reset," and "System Reset" in this chapter.)

Some channels can tolerate an absence of data transfer during a burst-mode operation, such as when a long gap on tape is read, but for not more than approximately half a minute. An equipment malfunction may be recognized when the absence of data transfer exceeds this time.

2.3.5 Ending Sequence. An ending sequence may be initiated by either the I/O device or the channel. If the procedure is initiated by the I/O device, then ending

status is presented, signalling the end of an operation. The end of an operation may be signalled by using only one ending sequence, assuming that both channel-end and device-end status conditions occur together, or the execution of two ending sequences may be required, assuming that device-end status had not been generated at the time channel-end status was presented. If the sequence is initiated by the channel, the I/O device may still require time to reach the point where the proper ending status information is available, in which case a second ending sequence is necessary to complete the ending procedure.

One of three situations may exist that result in the initiation of the ending sequence (assuming selection is already obtained):

1. The channel recognizes the end of an operation before the I/O device reaches its ending point. In this situation, whenever the control unit requires service again, the control unit raises either 'service in' or 'data in.' The channel responds by raising 'command out,' indicating stop. The control unit drops 'service in' or 'data in,' respectively, and proceeds to its normal ending point without requesting further service. When the I/O device reaches the point at which it would normally send 'channel end,' the control unit places the ending status on 'bus in' and raises 'status in.' The channel responds by raising 'service out,' unless it is necessary to stack the status, in which case the channel responds with 'command out.'

2. The channel and the I/O device recognize the end of an operation simultaneously.

3. The I/O device recognizes the end of an operation before the channel reaches the end of an I/O operation.

For situations 2 and 3, if the control unit is currently connected and status information is available, the control unit places the ending status on 'bus in' and raises 'status in.' If the control unit is not currently connected when the end of the I/O operation is recognized by the I/O device, a control-unit-initiated sequence is performed, the control unit becomes connected to the I/O interface, and then the status is presented by executing an ending sequence.

If device end is not presented with channel end, device end is presented when it is available by executing a control-unit-initiated sequence.

2.3.6 Addressing. An eight-bit address byte (plus parity) is used over the I/O interface for direct addressing of attached I/O devices. A unique eight-bit I/O-device address is assigned to each I/O device at the time a control unit is installed.

2.3.6.1 Address Assignment. At the time of installation, control-unit and I/O-device addresses are assigned as follows:

1. I/O devices that do not share a control unit with other devices may be assigned any I/O device address in the range 0-255, provided the address is not recognized by any other control unit attached to that I/O interface. Logically, such I/O devices are not distinguishable from their control unit, and both are identified by the same address.

2. I/O devices sharing a control unit (for example, magnetic-tape units and disk-storage units) are assigned addresses within sets of contiguous numbers. The size of a set of contiguous numbers is equal to the maximum number of I/O devices that can share the control unit, or 16, whichever is smaller. Furthermore, the set of addresses starts with an address in which the number of rightmost zeros is at least equal to the number of bit positions required to specify the size of the set. The leftmost bit positions of an address within a set identify the control unit; the rightmost bit positions designate the I/O device on the control unit.

3. Control units designed to accommodate more than 16 I/O devices may be assigned nonsequential sets of addresses. Each set consists of 16 addresses, or the number required to make the total number of assigned addresses equal to the maximum number of I/O devices attachable to the control unit, whichever is smaller. The addressing facilities are added in increments of a set so that the number of I/O-device addresses assigned to a control unit does not exceed the number of I/O devices attached by more than 15. For example, if a communication-type control unit has a designed capacity of 56 direct-access paths and if only 40 I/O devices are to be installed at the time of installation, 48 I/O device addresses can be assigned. However, if the full capacity of 56 I/O devices is to be installed, exactly 56 addresses are assigned.

I/O devices accessible through more than one path in the same system may have a different control-unit address for each path of communications. For sets of I/O devices connected to two or more control units, the portion of the address identifying the I/O device on the control unit is fixed and does not depend on the path of communications.

Except as indicated in the preceding rules, the assignment of control-unit and I/O-device addresses is arbitrary. The assignment is made at the time of installation, and the addresses normally remain fixed thereafter.

2.3.6.2 Address Decoding. Control units recognize an I/O-device address that meets the following conditions:

1. The address has correct parity.

2. The address is assigned to the control unit.

The control unit does not respond to any address outside its assigned set or sets. For example, if a control unit is designed to control I/O devices that have only bits 0000-1001 in the rightmost positions of the device address, the control unit does not recognize addresses that have 1010-1111 in these bit positions. If no control unit responds to an address, 'select out' is propagated through all control units and sent back to the channel as 'select in.' This may occur because:

1. An I/O-device address is not installed, or

2. An I/O-device address has been partitioned out of the system by the program, operator, or customer engineer.

The control unit must respond to those addresses in the set for which the corresponding I/O devices are:

1. Ready.

2. Not ready but which can be made ready by means of an ordinary manual intervention. The not-ready state of an I/O device is indicated by unit-check status and sense data that specifies intervention is required. (See "Intervention Required" in this chapter.)

The control unit may respond to all addresses in the assigned set, regardless of whether the device associated with the address is installed. If a control unit responds to an address for which no I/O device is installed, the unit-check status is set, and the appropriate sense information is made available.

The portion of the address decoder which identifies the control unit (item 2 under "Address Assignment" in this chapter) can be set at the time of installation for any bit combination.

Control units that are designed to attach only a single I/O device must decode all eight bits of the address byte. The address decoder can be set to any bit combination at the time of installation. (See item 1 under "Address Assignment" in section 2.3.6.1.)

2.4 Commands. When 'command out' is up during the initial-selection sequence, the byte of information on 'bus out' describes a command. The command specifies to the I/O device the operation to be performed.

The rightmost bit positions indicate the type of operation; the leftmost bit positions indicate a modification code which expands the meaning of the basic operation that is to be performed.

The modifier codes and the operations performed when they are decoded are specified in the vendor's manual for the I/O device.

The command byte on the I/O interface is defined as follows, where "M" is the modifier bit and "P" the parity bit:

	Bit Position								
	P	0	1	2	3	4	5	6	7
Test	P	M	M	M	M	0	0	0	0
Test I/O	1	0	0	0	0	0	0	0	0
Reserved	0	0	0	0	1	0	0	0	0
				through					
Reserved	1	1	1	1	1	0	0	0	0
Sense	P	M	M	M	M	0	1	0	0
Basic sense	0	0	0	0	0	0	1	0	0
Sense ID	1	1	1	1	0	0	1	0	0
Reserved	P	M	M	M	M	1	0	0	0
Read backward	P	M	M	M	M	1	1	0	0
Write	P	M	M	M	M	M	M	0	1
Read	P	M	M	M	M	M	M	1	0
Basic read	0	0	0	0	0	0	0	1	0
Control	P	M	M	M	M	M	M	1	1
No-operation	1	0	0	0	0	0	0	1	1

The basic sense, test I/O, no-operation and sense ID are executed on all I/O devices.

2.4.1 Basic Operations. The I/O operation to be executed over the I/O interface is determined by the eight-bit command issued to the I/O device during the initial-selection sequence.

The basic operations are specified by these commands: read, read backward, write, control, sense, and test.

A command with invalid parity is not recognized and therefore not executed.

2.4.1.1 Immediate Operation. Some commands cause the I/O device to signal channel end as initial status during an initial-selection sequence. An I/O operation causing channel end to be signalled as initial status (with the exception of test I/O) is called an immediate operation.

An immediate operation is performed when the execution of that command meets the following requirements:

1. Execution requires no more information than that in the command byte; that is, no data bytes are transferred.

2. Channel-end status is presented as initial-status. Device-end status may also accompany channel end.

Any command (with the exception of test I/O) may be executed as an immediate operation.

Notes:
1. The inadvertent use of special diagnostic commands, which by intention permit errors to occur on the I/O interface or which introduce the possibility that subsequent commands may be executed erroneously, must be prevented by some form of interlock.
2. An error condition may be recognized by the channel and the I/O operation terminated when 256 or more chained commands are executed with a single I/O device and none of the commands result in the transfer of any data.

2.4.1.2 Read. The read command initiates execution of data transfer from the control unit to the channel. The bytes of data within a block are provided in the same sequence as those written by the write command.

A read command with all modifier bits set to zeros is a basic read command. This command is also used as an initial-program-loading (IPL) read command by those devices that provide the IPL function. To perform the initial-program-loading read, the command must be:

1. The first command sent to the I/O device following a system reset.

2. Sent no sooner than 1 millisecond following system reset.

2.4.1.3 Read Backward. The read-backward command initiates an I/O operation in the same manner as the read command, except that bytes of data within a block are sent to the channel in an order which is the reverse of that used in writing. The control unit may be designed to cause mechanical motion in the I/O device in a direction opposite of that for a read command, or it may be designed to operate the device as it would for a read command.

Unless otherwise noted, any description that applies to read also applies to read backward.

2.4.1.4 Write. The write command performs the same sequence of signals over the I/O interface as for a read operation. For a write command, however, the data is sent from the channel to the control unit instead of from the control unit to the channel.

2.4.1.5 Control. The control command is similar to the write command, except that the command modifier bits received by the control unit are decoded to determine which of several possible functions are to be performed. One function performed may be second-level addressing, which may require several bytes of data to complete the control operation. When the particular control function can be channel-end, status may be presented during the initial-selection sequence.

The byte rate for the bytes transferred during an operation specified by a control command normally need not be faster than the normal read or write rate for the same I/O device. Exceptions are specified in the vendor's manual for the I/O device.

A control command with all-zero modifier bits performs no operation at the I/O device, except that it satisfies any previously indicated chaining operations and allows certain I/O devices to wait for conditions of checking (or any synchronizing indications) before releasing the channel. This variation of the control command is a no-operation control.

2.4.1.6 Sense. The sense command is similar to a read command, except that the data is obtained from sense indicators rather than from a record source.

The basic sense command (modifier bits set to zero) initiates a sense operation on all I/O devices and causes the retrieval of up to 32 bytes of data. The basic sense command does not initiate any operation other than the sensing of sense indicators. The basic sense command sent to an available control unit is accepted even though the addressed I/O device is in a not-ready state. (See section 2.7 "Sense Information.") If the control unit detects an error during the sense operation, unit check is sent with the channel-end status condition.

The purpose of the basic sense command is to provide data detailed enough to ascertain the actual state of the device and unusual conditions associated with the execution of the I/O operation during which the error was detected.

The I/O devices that can provide special diagnostic sense information or that can be instructed to perform other special functions by means of the sense command may call for modifier bits to be defined for control of the function. The special sense operation may be initiated by unique combinations of modifier bits, or a group of codes may specify the sense function. Any remaining sense command codes may be considered invalid or may cause the same action as the basic sense command, depending on the particular I/O device.

The sense-ID command (modifier bits set to E hex) does not initiate any operations other than sensing the type/model number.[4] If the control unit or device is available and not busy, then the sense-ID command is executed, with up to seven bytes of data transferred. Basic sense data may or may not be reset as a result of executing the sense-ID command.

The sense bytes sent in response to sense ID are:

Bytes	Contents
0	FF hex
1,2	Control-unit type number
3	Control-unit model number
4,5	I/O-device type number
6	I/O-device model number

All unused sense bytes must be set to zeros.

[4]Many older units may properly execute the sense-ID command, may execute the command as a basic sense command, or may reject the command with unit-check status.

In bytes 1 and 2, the four-decimal-digit number corresponds directly with the unit type number on the tag attached to the unit.

Byte 3 is all zeros if the number is not applicable.

In bytes 4 and 5, the four-decimal-digit number corresponds directly with the I/O-device type number on the tag attached to the I/O device.

Byte 6 is all zeros if the number is not applicable.

Whenever a control unit is not separately addressable from the attached I/O device or devices, the response to the sense-ID command is a concatenation of the control-unit type number and the I/O-device type number.

If a control unit can be addressed separately from the attached I/O device or I/O devices, then the response to the sense-ID command depends on the unit addressed. If a control unit is addressed, the response to the sense-ID command is the control-unit type number, with normal ending status presented after byte 3. If the I/O device is addressed, the response to the sense-ID command is the I/O-device type number in bytes 1, 2, and 3, with normal ending status occurring after byte 3.

2.4.1.7 Test. The test-I/O command (modifier bits set to zeros) retrieves from the addressed I/O-device path any stacked or pending status and does not initiate an operation. If no stacked or pending status is encountered along the I/O path being tested, a zero status byte for the selected I/O device is presented to the channel.

With respect to test I/O, the busy condition has a special meaning. (See section 2.6.6 "Busy.")

Note: Any status presented during execution of a test-I/O command may be stacked.

2.5 Sequence Controls. Sequence controls are used by the channel to control execution of the sequences that are performed between the channel and control unit. Each sequence control uses a special signalling convention over the I/O interface, and each one has a particular meaning to the control unit or I/O device.

2.5.1 Proceed. Whenever 'command out' rises in response to the rise of 'address in' during a control-unit-initiated sequence, it means proceed.

Proceed indicates to the I/O device that it should continue the normal servicing request on the interface.

2.5.2 Stop.

Note: When the data-streaming feature is used, this "Stop" section does not apply. (See section 2.9.5 "Data-Streaming Feature.")

The channel uses the stop sequence control when data is being transferred and the channel recognizes that the currently executing I/O operation should be ended. 'Stop' is indicated by raising 'command out' in response to the rise of 'service in' or 'data in.'

On receipt of the stop signal, the I/O device proceeds to its normal ending point without sending any further 'service in' or 'data in' signals to the channel. The I/O device remains busy until the ending status is available, is presented to the channel, and is accepted by the channel.

During a data-transfer sequence, 'command out' is transmitted in response to the first 'service in' or 'data in' that is provided after the channel determines that the current operation is to be ended. If 'select out' is down or goes down after this sequence, 'operational in' drops on force-burst-mode-type operations on I/O devices that cannot meet the timeout requirements indicated in "Interface Timeout Considerations." Also, burst-mode I/O devices that have relatively long times between stop and ending status and have no time-dependent chaining requirements drop 'operational in' at this time. (See the appropriate Operational Specification.)

2.5.3 Stack Status. The channel uses the stack sequence control when conditions preclude acceptance of status from the control unit. Stack status is indicated by the rise of 'command out' in response to the rise of 'status in.' The stack-status signal causes retention of status information at the control unit or I/O device until that status is accepted during a subsequent sequence. When 'stack status' occurs, the control unit disconnects from the interface after 'select out' is down. 'Command out' remains up until 'operational in' falls. Any attempt to perform a control-unit-initiated sequence in order to present the status again is under control of 'suppress out.' (See "Suppress Status" in this chapter.) Any status (except zero status presented in response to a command other than test I/O) presented by a control unit in any interface sequence (except the short-busy sequence) may be stacked.

If the channel signals stack status to a control unit as status is being presented, command chaining, if any, is not indicated when that status is subsequently accepted by the channel.

2.5.4 Suppress Data. For control units whose rate of data transfer can be adjusted without causing an overrun condition, the channel may use the suppress-data sequence control. The suppression of data occurs according to the following rules. (Operations with completely buffered I/O devices or start-stop devices and the transfer of data for the basic sense command may fall into this category.)

1. Except when the data-streaming feature is used, 'suppress out' is ignored for the first data byte following any selection sequence unless the data transfer is contiguous with initial selection, that is, unless there is no disconnection and reconnection between initial selection and the data-transfer sequence. (See section 2.9.5 "Data-Streaming Feature.")

2. To ensure recognition by the control unit, 'suppress out' must be up 250 nanoseconds before the rise of 'service in' or 'data in,' or at least 250 nanoseconds before 'service out' or 'data out' falls. (See section 2.9.4 "High-Speed-Transfer Feature.")

3. When 'suppress out' is up at the control unit and the operation is in burst mode (either because 'select out' is up or because the control unit is forcing burst mode), the control unit does not raise 'service in' for subsequent suppressible data.

2.5.5 Accept Data.

Note: When the data-streaming feature is used, this "Accept Data" section does not apply. (See section 2.9.5 "Data-Streaming Feature.")

The channel uses the accept-data sequence control during data transfer from the control unit to the channel. When 'service out' or 'data out' rises in response to 'service in' or 'data in,' respectively, during a read, read-backward, or sense operation, the channel signals acceptance of the information placed on 'bus in' by the control unit.

2.5.6 Data Ready.

Note: When the data-streaming feature is used, this "Data Ready" section does not apply. (See section 2.9.5 "Data-Streaming Feature.")

The channel uses the data-ready sequence control during data transfer from the channel to the control unit. When 'service out' or 'data out' rises in response to 'service in' or 'data in,' respectively, during a write or control operation, the channel signals that the requested information has been placed on 'bus out' and is ready for acceptance by the control unit.

2.5.7 Suppress Status. Whenever the channel is unable to immediately accept status, the suppress-status sequence control may be used. When 'suppress out' is raised, the control unit does not attempt a control-unit-initiated sequence to present suppressible status.

It is acceptable for a control unit to treat all status as being suppressible only after that status has been stacked. Alternatively, it is acceptable for a control unit to treat status as being suppressible if (1) that status contains channel end and interface disconnect has been received previously during the I/O operation, or (2) that status contains device end that ended an I/O operation when command chaining was not indicated at the time channel end was presented. Also, asynchronous status may be suppressible at the option of the particular control unit without being stacked.

'Suppress out' must be up at least 250 nanoseconds before 'select out' rises at the control unit to ensure suppression of status. 'Suppress out' suppresses the initiation of the control-unit-initiated sequence when the sequence is intended to present suppressible-type status. If 'suppress out' rises after a control-unit-initiated sequence has been started, the status sequence proceeds normally.

The relationship between 'request in' and 'suppress out' is described in "Request In" in Chapter 2, "Functional Specifications."

2.5.8 Accept Status. During presentation of status, the channel may use the accept-status sequence control. When 'service out' rises in response to 'status in,' the channel signals that the status placed on 'bus in' by the control unit has been acccepted. 'Service out' falls in response to the fall of 'status in.'

2.5.9 Command Chaining. During the execution of successive I/O operations, the channel uses the command chaining sequence control. Command chaining is indicated if 'suppress out' is up when 'service out' is raised in response to 'status in.' When channel end and device end are presented together and if command chaining is to occur, it is indicated when the status is accepted by the channel. When channel end and device end are not presented together and if command chaining is to occur, it is indicated when channel-end status is accepted and again when device-end status is accepted. When command chaining is indicated, the command chaining condition is set in the control unit.

Command chaining means that another initial-selection sequence is to occur for the I/O device in operation immediately following the presentation of device end, provided that no unusual conditions were encountered during execution of the current operation. The exact time at which the next command is presented depends on the channel and on the parameters of the system.

If command chaining was indicated during the presentation of channel end but is not indicated when device end is presented, command chaining is suppressed, and the command chaining condition is reset in the control unit.

When command chaining is indicated at the time device end is presented, this indication is valid until reselection is made or until 'suppress out' falls (minimum down-level time to ensure a break is 250 nanoseconds). To ensure that command chaining occurs, 'suppress out' remains up during the reselection at least until 'operational in' rises. If 'suppress out' drops before the reselection is made, the command chaining condition is reset in the control unit.

Reselection of any I/O device attached to a multidevice control unit that cannot have multiple operations on multiple devices in progress at one time resets the command chaining condition in the control unit. A multidevice control unit that can have multiple operations on multiple devices in progress at one time maintains the command chaining condition for each device. This type of control unit does not reset the command chaining condition for devices other than for the one being selected.

When the channel accepts status from a control unit that contains unit check, except when command retry is signalled, unit exception, busy, control-unit end, or attention, the channel terminates command chaining in the channel, and the control unit resets the command chaining condition for the device whose address was used to present that status.

The channel terminates command chaining in the channel, and the command chaining condition is reset in the control unit whenever interface disconnect, selective reset, or system reset is signalled during the execution of an I/O operation.

Depending on the particular I/O device, operation, and configuration, the command chaining indication requires certain functional control that depends on the individual control unit.

If command chaining is indicated when channel end but not device end is presented from a multidevice control unit and if the control unit cannot have multiple

operations on multiple devices in progress at one time, the I/O device that presents channel end is the next I/O device from that control unit to present device-end status, unless the control unit is addressed in the meantime on the same I/O interface.

If command chaining is indicated when device end is presented, the control unit ensures that the path to the I/O device remains available until the next initial-selection sequence is initiated that immediately follows the acceptance of device end or until command chaining is no longer indicated. Furthermore, unless command chaining is being canceled by the channel, the immediately following sequence is a reselection of the I/O device presenting the device end.

If command chaining is indicated on an I/O device shared between more than one control unit or channel, the I/O device remains available until the initial-selection sequence is initiated that immediately follows the acceptance of device end or until command chaining is no longer indicated.

To ensure recognition of command chaining by the control unit, 'suppress out' is up at least 250 nanoseconds before 'service out' rises in response to 'status in' and does not fall before 'status in.' If command chaining is not to be indicated, 'suppress out' is down at least 250 nanoseconds before the rise of 'service out' and does not rise before the fall of 'status in.'

2.5.10 Interface Disconnect. During the execution of an I/O operation, the interface-disconnect sequence control may be used by the channel to signal the control unit to end execution.

If 'select out' ('hold out') is down and 'address out' rises or if 'address out' is up and 'select out' ('hold out') falls, the presently connected control unit drops 'operational in,' thus disconnecting from the interface. Mechanical motion in process continues to a normal stopping point. Status information is generated and presented to the channel when appropriate. 'Address out,' in this case, may be up concurrently with another outbound tag line. Interface disconnect is signalled when 'address out' is up and 'select out' ('hold out') is down. To ensure recognition of interface disconnect, it must occur at least 250 nanoseconds before the rise of the outbound tag that completes any signal sequence.

'Operational in' must drop within 6 microseconds after the disconnect indication is received. When 'operational in' drops, the channel may drop 'address out' to complete the interface-disconnect sequence. 'Address out' must be down for at least 250 nanoseconds before a new initial-selection sequence is attempted.

The control unit responds to the interface disconnect signal by removing all signals (with the possible exception of 'request in' and 'metering in') from the I/O interface. On an input operation, data on 'bus in' need not be valid after the rise of 'address out.' On an output operation, data on 'bus out' must be valid until the fall of either 'service in,' 'data in,' or 'operational in.' (See section 2.9.5 "Data-Streaming Feature.") When the control unit reaches the normal ending point, it attempts to obtain selection on the interface to present any generated status to the channel.

To ensure recognition of interface disconnect during an initial selection sequence, interface disconnect should not be signalled in the interval during or subsequent to the rise of 'command out' and prior to or during the rise of 'service out' in response to 'status in.' If interface disconnect is signalled during this interval, it is unpredictable whether the command is executed and whether ending status should be sent by the control unit.

Any abnormal I/O device operation should be indicated by unit check in the status, and the sense information should provide additional details concerning the operation. (See "Unit Check" in this chapter.) The control unit does not generate any status solely as a result of an 'interface disconnect.'

The I/O-device path remains busy after it receives an 'interface disconnect,' while performing an operation, until device-end status is accepted by the channel. If 'interface disconnect' is received when the I/O device is not busy, no status is generated, and the I/O device is not made busy.

Note: Except when the data-streaming feature is used, if 'address out' is up concurrently with another out tag, the information on 'bus out' remains valid until the associated in tag drops or until 'operational in' drops. (See section 2.9.5 "Data-Streaming Feature.")

2.5.11 Selective Reset. The selective-reset sequence control is used when a malfunction is detected at the channel or when a control unit signals 'disconnect in.' (See section 2.9.2 "I/O-Error-Alert Feature.")

Selective reset is indicated whenever 'suppress out' is up and 'operational out' drops. This condition causes 'operational in' to fall and causes the particular I/O device in operation and its status to be reset. The operation in process proceeds to a normal stopping point, if applicable, with no further data transfer. (See section 2.9.5 "Data-Streaming Feature.") The I/O device operating over the interface is the only device that is reset, even on multidevice control units. The particular I/O-device path is in a busy state throughout this procedure.

If selective reset is to be recognized by the control unit, 'suppress out' rises at least 250 nanoseconds before 'operational out' drops and remains up until at least 250 nanoseconds after 'operational out' rises. 'Operational out' stays down until 'operational in' falls or for at least 6 microseconds, whichever is greater, for the selective reset to be effective.

The ready or not-ready state of the control unit is generally not changed by a selective reset. When, however, the enable/disable or online/offline switch was changed before the reset but had not become effective because of the required inhibiting conditions, the ready or not-ready-state may change if the reset clears those inhibiting conditions.

'Device end' may be returned after a selective reset has been signalled. The interpretation of the selective reset is defined in the vendor's manual for the I/O device.

2.5.12 System Reset. The system-reset sequence control is used to reset all control units and I/O devices that are online. (See section 2.8.4 "Online/Offline.") System reset is indicated whenever 'operational out' and 'suppress out' are down concurrently and the I/O device is in the online mode. This condition causes 'operational in' to fall and causes all control units and their attached I/O devices, along with their status, to be reset. The control units are in a busy state for the duration of their reset procedure. 'System reset' can prepare an I/O device for an initial-program-loading sequence.

The ready or not-ready state of the control unit is generally not changed by a system reset. When, however, the enable/disable or online/offline switch was changed before the reset but is not yet effective because of required inhibiting conditions, the ready or not-ready state may change if the reset clears those inhibiting conditions.

To ensure a proper reset, 'operational out' and 'suppress out' are down concurrently for at least 6 microseconds.

The interpretation of the system-reset signal by the I/O device is defined in the operational specification manual for that I/O device.

2.6 Status Information. When 'status in' is up, the information that appears on 'bus in' is the status byte. The conditions reported in the status byte are the status conditions.

The status pertains to the I/O device or control unit whose address appeared on 'bus in' (with 'address in') during the control-unit-initiated sequence or the initial-selection sequence. In the case of the short-busy sequence, when no 'address in' occurs, it is assumed that the status pertains to the addressed I/O device or control unit.

Note: Unless otherwise stated, the information in this standard pertains to control units attached to only one channel interface.

2.6.1 Status Byte. The status byte has the following format:

Bit Position	Designation
P	Parity
0	Attention
1	Status modifier
2	Control-unit end
3	Busy
4	Channel end
5	Device end
6	Unit check
7	Unit exception

The status byte is transmitted to the channel:

1. During the initial-selection sequence.

2. To present the channel-end status at the termination of data transfer.

3. To present the device-end status and any associated conditions to the channel. The I/O device remains busy during an operation until the channel accepts the device-end status.

4. To present control-unit-end or device-end status, which signals that the control unit or device that was previously busy and interrogated while busy is now free.

5. To present any previously stacked status when allowed to do so.

6. To present status describing asynchronous conditions that are recognized by either a control unit or an I/O device and that are unrelated to previously executed I/O operations. One of these conditions is described by attention, which is normally generated by console or communication devices. Another condition, which is described by device end, unit exception, and attention, is generated when the corresponding device goes from the not-ready to the ready state. These status conditions are handled in the same way as any other status information presented during a control-unit-initiated sequence and are subject to the same rules regarding presentation to the channel and stacking. (See section 2.6.8 "Device End" and section 2.3.3 "Control-Unit-Initiated Sequence.")

Once accepted by the channel, any given status byte is reset and is not presented again.

2.6.2 Unit-Status Conditions. The following status conditions are detected by the I/O device or control unit and are indicated to the channel over the I/O interface. The causes of these conditions for each type of device and the timing in presenting these conditions are specified in the vendor's manual for the I/O device. When such conditions have been recognized, status is generated at the I/O device or control unit and is maintained until it is accepted by the channel or is reset. Status that has been generated but not presented to the channel is called "pending status." Status that has been presented to the channel but has not been accepted and is currently being held at the control unit or I/O device is called "stacked status."

When the device is accessible from more than one channel, status resulting from channel-initiated operations is signalled to the channel that initiated the associated I/O operation. The handling of conditions not associated with I/O operations, such as attention, unit exception, and device end because of a not-ready-to-ready transition, depends on the type of device and condition and is specified in the operational specification for the I/O device. (See section 2.6.8 "Device End.")

Note: Control units and I/O devices should provide interlocks so that status is not lost, hidden, or included with other status when the result would cause the program to misinterpret the original meaning and intent of the status.

2.6.3 Attention. The attention condition is generated when some asynchronous condition occurs in the I/O device. The condition may be accompanied by other status. Attention is not associated with the initiation, execution, or termination of any I/O operation.

If a condition at the I/O device or control unit has caused the presentation of attention status while an I/O operation is in execution, command chaining, if any, is no longer indicated, and the operation is ended. If command chaining is indicated when channel-end status is presented, command chaining is no longer indicated when device-end status is presented and accompanied by attention. When the attention condition is indicated, the handling and presentation of the condition to the channel depends on the type of I/O device. An I/O device shared between more than one channel path presents the attention status to all channel paths as defined in the vendor's manual for that I/O device. Depending on the I/O-device application, attention may or may not be presented until command chaining is no longer indicated.

Attention is accompanied by device end and unit exception when a not-ready-to-ready-state transition is signalled. (See section 2.6.8 "Device End.")

2.6.4 Status Modifier. Status modifier is used by control units in five cases:

1. Control units that cannnot provide current status in response to test I/O present the status-modifier bit alone during the initial-selection procedure.

2. Busy control units present the status-modifier bit with the busy bit during the initial-selection sequence or short-busy sequence to differentiate between a busy control unit and a busy device.

3. Control units designed to recognize special ending conditions (search equal on a disk) present the status-modifier bit with device end when the special condition occurs.

4. Control units that use the command-retry feature present status modifier with unit check and channel end--with or prior to device end--when requesting that the channel initiate the retry procedure.

5. Control units designed to recognize special conditions that must be brought to the attention of the program present status modifier along with the other status indications in order to modify the meaning of the status.

In the first case, provision is made for certain types of control units that cannot provide current status on a demand basis as required by test I/O. Presentation of the status-modifier bit in response to test I/O indicates that the control unit cannot execute the command, and any existing status remains unchanged and unavailable to a test I/O. This type of control unit provides status only on a control-unit-initiated selection sequence.

In the second case, provision is made for indicating that a busy condition pertains to a control unit and not necessarily to the addressed I/O device. (See section 2.6.6 "Busy" for a description of control-unit-busy status conditions.)

In the third case, provision is made for control units designed to recognize special ending or synchronizing conditions. If the special condition occurs, the

status-modifier bit with the device-end bit is presented during the status presentation. When status-modifier and device-end bits are present in the status, it indicates that the normal sequence of commands must be modified.

In the fourth case, provision is made for control units to request a retry of the execution of the current command. The situations encountered during which retry is requested are specified in the operational specifications for the device, or in the vendor's documentation for the device.

In the fifth case, provision is made for control units designed to recognize special conditions that are unrelated to the execution of an I/O operation. These conditions must be brought to the attention of the program. Status modifier is used to modify the meaning of the other accompanying status indications. The meaning of the status combinations is specified in the operational specifications for the device.

2.6.5 Control-Unit End. Only control units that can indicate a control-unit-busy condition can indicate a control-unit-end condition. The control-unit-end condition occurs for one or both of the following:

1. The control unit was interrogated while it was in the busy state. Interrogated in the busy state means that, during an initial-selection sequence, conditions at the control unit precluded execution of an operation with an I/O device on the control unit and that the control unit responded with busy and status modifier in the unit-status byte. (See section 2.6.4 "Status Modifier.")

2. The control unit detected an unusual condition while busy, but after channel end was accepted by the channel. Indication of the unusual condition accompanies control-unit end.

If the control unit remains busy executing an operation after signalling channel end but is not interrogated and does not detect an unusual condition, control-unit end is not generated.

The I/O-device address associated with control-unit end is determined as follows:

1. The address of the selected device is used if control-unit end is to be presented with channel end and/or device end.

2. If control-unit end is generated without channel end or device end and the status is presented during a control-unit-initiated sequence, the I/O-device address to be used when presenting this status is any legitimate address associated with the control unit. (A legitimate address is any address that is not associated with an I/O device currently executing an I/O operation and that the control unit is capable of recognizing, regardless of whether or not the I/O device is actually attached.)

3. If control-unit end is to be presented during an initial-selection sequence, the I/O-device address is the same as the address issued with 'address out.' The sequence may be executed as either the initial-selection sequence or the short-busy sequence.

The control-unit-end condition can be signalled with channel end, or any time after channel end, and may be accompanied by other status bits. When control-unit end is signalled with a control-unit-initiated sequence in the absence of any preceding data transfer or other status conditions, the status may be associated with any address assigned to the control unit. For control units attaching more than a single I/O device, a pending control-unit end for one I/O device does not necessarily preclude initiation of new operations with other attached devices. Whether the control unit allows initiation of other operations is at the option of the control unit. Control-unit end causes command chaining to be suppressed.

2.6.5.1 Temporary Control-Unit Busy. When the busy state of the control unit is temporary, control-unit end can be included with busy and status modifier in response to interrogation, even though the control unit is not free. The busy condition may be considered temporary if it lasts less than approximately 2 milliseconds.

Note: When a control unit has signalled the temporary control-unit-busy condition, the channel may attempt an initial-selection sequence immediately following acceptance of the status.

2.6.6 Busy. The busy indication occurs only during an initial-selection sequence or short-busy sequence and means that conditions existing at the I/O device or control unit preclude execution of the intended I/O operation because one of these four situations exists:

1. An I/O operation initiated during a previous initial-selection sequence is being executed.

2. Stacked or pending status conditions exist (except as noted later in this section).

3. The control unit is shared by channels or I/O devices, or an I/O device is shared by control units and the shared facility is not available.

4. A self-initiated function (for example, microdiagnostics or data movement internal to the I/O device) is being performed.

An I/O operation is being executed from the time initial status is accepted until device end is accepted. Status conditions for the addressed I/O device, if any, accompany the busy indication.

If the busy condition applies to a control-unit function, busy is accompanied by status modifier.

Busy is indicated to a test-I/O command under the same conditions that busy is indicated to other commands; however, busy is not indicated to a test-I/O command if there is an available path to the I/O device and there is stacked or pending status for that I/O device.

The busy condition causes command chaining to be suppressed.

Note: For many recent I/O devices if zero status has previously been stacked at a device and if a command other than test I/O is received by the same device, the zero status is discarded by the control unit or device, and execution of the command proceeds normally.

2.6.7 Channel End. Channel end is caused by the completion of the portion of an I/O operation involving execution of a command; this execution may also include the transfer of data between the I/O device and the channel.

Each I/O operation causes only one channel-end indication to be generated. The channel-end condition is generated only when the command is accepted, that is, the initial status byte for the operation contained either all zeros or channel end, as in the execution of an immediate operation. If the initial-status byte contained all zeros, the next status byte presented by the device includes channel end. If channel end is not included in this status byte, an error condition may be recognized by the channel and the I/O operation, if any, terminated.

During an I/O operation, the exact time that channel end is generated depends on the operation and the type of I/O device. For operations such as writing, some I/O devices generate the channel-end condition when the data has been written. On other I/O devices that later verify the writing, channel end may be delayed until verification is performed, depending on the I/O device. On I/O devices equipped with buffers, the channel-end condition may occur on completion of data transfer between the channel and the buffer. During execution of control commands that involve data transfer, channel end is usually generated after the data is transferred to the control unit, although for some I/O operations, channel end may be delayed until the operation is completed. Operations that do not cause data to be transferred, as in the case of immediate operations, can provide the channel-end condition during the initial-selection sequence.

2.6.8 Device End. Device end is presented to the channel (1) when the completion of an I/O operation is signalled by the I/O device, (2) when the I/O device, having previously responded busy, signals that a change from the busy to the not-busy state has occurred, (3) when the I/O device signals that a change from the not-ready to the ready state has occurred, and (4) when the control unit or I/O device signals that an asynchronous condition has been recognized.

Each I/O operation cuases only one device-end condition. In this situation, the device-end condition is not generated unless the command is accepted.

The device-end condition associated with an I/O operation is generated either simultaneously with the channel-end condition or later. For data-transfer operations on some I/O devices, the operation is complete at the time channel end is generated, and both device end and channel end occur together. The time at which device end is presented depends upon the I/O device type and the kind of command executed. For most I/O devices, device end is presented when the I/O operation is completed. In some cases, for reasons of performance, device end is presented before the I/O operation has actually been completed at the I/O device. However, in all caes, when device end is presented, the I/O device is available for an initial-selection sequence if command chaining was indicated when the device end was accepted by the channel. During execution of control commands, the

device-end condition may be generated at the time channel end is generated, or later.

When command chaining is being performed, the channel normally makes available to the program only the device end of the last operation of the chain. In situations where the I/O operation has been completed and status of only channel end has been accepted by the channel, the next status byte presented by the I/O device contains either device end or device end and status modifier; otherwise, command chaining, if any, is not indicated. A device-end signal received by the channel in the absence of any unusual conditions causes the channel to initiate the next command-chained I/O operation, if any. Device-end status, when received by the channel in the presence of unusual conditions, causes the channel to terminate command chaining, if any, and report the unusual condition to the program. If, during command chaining, an unusual condition is recognized subsequent to device end having been accepted by the channel and before acceptance of the next command by the control unit, command chaining does not occur, and the device-end status previously presented is not made available to the program.

If an I/O device previously responded busy, device end is signalled on the path over which the initial-selection sequence was initiated when the I/O device becomes not busy. In this situation, device end is signalled only once, independent of the number of times the I/O device responded busy.

Device end is accompanied by attention and unit excepiton when a state change from not-ready to ready is signalled.[5] An I/O device is considered to be not ready when operator intervention is required in making the I/O device ready. A not-ready condition can occur, for example, because of any of the following:

1. An unloaded condition in a magnetic-tape unit.
2. Card equipment out of cards or stacker full.
3. Printer out of paper.
4. Error conditions that need operator intervention.
5. An I/O device having changed from the enabled to the disabled state.

An I/O device shared between more than one channel and for which a not-ready-to-ready-state transition has occurred must present the appropriate indication[5] to all attached channels as defined in the operational specification for the I/O device.

Device end is accompanied by other status when conditions are recognized that are unrelated to the execution of an I/O operation. These conditions are brought to the attention of the program and are presented to the channel as they occur.

2.6.9 Unit Check. Unit check indicates that the I/O device or control unit has detected an unusual condition that is detailed by information available to a basic sense command. The occurrence of unit check may indicate that a programming error or an equipment error has been detected, that the not-ready state of the

[5]Many older I/O devices present status of only device end or device end, attention, and unit exception when signalling a not-ready-to-ready-state transition.

device has affected the execution of the command, or that an exceptional condition other than the one identified by unit exception has occurred. The unit-check bit provides a summary indication of the conditions identified by sense data. (See also section 2.9.3 "Command-Retry Feature.")

An error condition causes the unit-check indication when the error condition occurs during execution of a command or during some activity associated with an I/O operation. Unless the error condition pertains to the activity initiated by a command or is of significance to the program, the condition does not cause the program to be alerted after device end has been cleared; a malfunction may, however, cause the device to become not ready.

Unit check is indicated when the existence of the not-ready state precludes a satisfactory execution of the command or when the command, by its nature, tests the state of the I/O device. When no status is pending for the addressed I/O device at the control unit, the control unit signals unit check when test I/O or the no-operation control command is issued to a not-ready I/O device. In the case of a no-operation, the command is not accepted, and channel end and device end do not accompany unit check.

Unless the command is designed to cause unit check, such as a command to rewind and unload magnetic tape, unit check is not indicated if the command is properly executed, even though the I/O device has become not ready during or because of the operation. Similarly, unit check is not indicated if the command can be executed with the I/O device not ready. A console could, for example, accept and execute the alarm-control command when the printer is not ready. Selection of an I/O device in the not-ready state does not cause a unit-check indication when the sense command is issued or when a status condition is pending for the addressed device at the control unit.

If, during the initial-selection sequence, the I/O device detects that the command cannot be executed before initial status is presented, unit check is presented to the channel and it appears without channel end, control-unit end, or device end. Such unit status indicates that no action has been taken at the device in response to the command. If the condition that precludes proper execution of the operation occurs after the command has been accepted, unit check is accompanied by channel end, control-unit end, or device end, depending on when the condition was detected. Errors detected after device end is cleared may be indicated by signalling unit check with attention, unit check with control-unit end, or unit check with device end.

Unit-check status presented either in the absence of or accompanied by other status indicates only that sense information is available to the basic sense command. Presentation of either channel end and unit check or channel end, device end, and unit check does not provide any indication as to the kind of conditions encountered by the control unit, the state of the I/O device, or whether execution of the operation has been started. Instead, descriptions of these conditions or states are provided in the sense information.

When unit check appears with channel end and without device end, the sense data and an available device path must be preserved until after the device end and the sense data are accepted or reset.

Errors, such as invalid command code or invalid command-code parity, do not cuase unit check when the device is working or contains a pending interruption condition at the time of selection. Under these circumstances, the I/O device responds by providing the busy bit and indicating any pending status. The command-code validity is not indicated.

Termination of an operation with the unit-check indication causes command chaining to be suppressed.

Note: If an I/O device becomes not ready on completion of a command, the ending status can be cleared by test I/O without generation of unit check because of the not-ready state. Any subsequent test I/O issued to the I/O device causes a unit-check indication.

2.6.10 Unit Exception. Unit exception means that the I/O device detected an unusual condition that needs to be reported to the program. During execution of an I/O operation, unit exception has only one meaning for any particular command and type of I/O device. A sense operation is not required as a response to the acceptance of a unit-exception condition.

A unit-exception condition may be generated when the I/O device is executing an I/O operation, or when the device is involved with some activity associated with an I/O operation and the condition is of immediate significance. If a device detects a unit-exception condition during the initial-selection sequence, unit exception is presented to the channel without channel end, control-unit end, or device end. Such unit status indicates that no action has been taken at the device in response to the command. If the condition that precludes normal execution of the operation occurs after the command has been accepted, unit exception is accompanied by channel end, control-unit end, or device end, depending on when the condition is detected. Any unusual condition associated with an operation but detected after device end is cleared, is indicated by signalling unit exception with attention.

The unit-exception condition causes command chaining to be suppressed.

Unit exception is accompanied by device end and attention when a not-ready-to-ready-state transition is signalled. (See section 2.6.8 "Device End.")

2.7 Sense Information. Data transferred during a sense operation provides information concerning unusual conditions detected in a previous I/O operation and concerning the actual state of the I/O device. Information provided by the basic sense operation is more detailed than that supplied by the unit-status byte, and may describe reasons for the unit-check indication. It may also indicate, for example, that the I/O device is in the not-ready state, that a tape drive is in the file-protected state, or that magnetic tape is positioned beyond the end-of-tape marker.

Error information used for the recovery of the I/O operation, if any, normally is provided in the first byte of sense data. (See section 2.7.1 "Sense Byte.") All I/O devices provide at least the first sense byte and may transfer up to 31 additional sense bytes during execution of the basic sense command. (See section 2.4.1.6 "Sense.") The amount and the meaning of the additional sense data are peculiar to the type of device, and are specified in the vendor's manual for the I/O device.

A device which normally operates in byte-multiplex mode and transfers a single byte during data transfer should transmit the maximum number of sense bytes possible during the 32-microsecond byte-multiplex timeout limitation.

The sense information that pertains to a previous I/O operation or other unit action at an I/O device may be reset any time after the completion of the basic sense command addressed to that I/O device. Except for the test I/O and no-operation commands, any other command addressed to the control unit may be allowed to reset the sense information, provided that the busy bit is not included in the initial status. The sense information may also be changed as a result of asynchronous actions, such as when not-ready-to-ready status is generated.

Sense information that results from more than one action at the unit is not ORed when this condition would cause the program to misinterpret the original meaning and intent of the sense information. When a group of sense indicators is shared with different devices, the residual control-unit sense data that pertains to a previous command addressed to the control unit may be reset if the I/O device addressed is different from the I/O device which generated the sense data.

A command code with invalid parity causes the sense information to be replaced only if unit check is turned on as a result of the invalid parity.

2.7.1 Sense Byte. The first six bits of the first sense data byte (sense byte 0) are common to all I/O devices. The six bits are independent of each other and, when set to ones, designate the following:

Bit	Description
0	Command reject
1	Intervention required
2	Bus-out check
3	Equipment check
4	Data check
5	Overrun

2.7.2 Sense Conditions.

2.7.2.1 Command Reject. The command-reject condition occurs when the I/O device has detected a programming error. The command-reject condition is generated when a command is received which (1) the I/O device is not designed to execute (such as a read backward issued to a direct-access storage device) or (2) the I/O device cannot execute because of its present state (such as a write issued to a file-protected tape unit). In case 2, the program may have required use of an uninstalled optional feature or may have specified invalid control data. Command reject is also indicated when an invalid sequence of commands is recognized (such as a write to a direct-access storage device without the data block having previously been designated).

2.7.2.2 Intervention Required. The intervention-required condition occurs when the command could not be executed because of a condition that requires intervention at the I/O device. it may indicate conditions such as an empty hopper in a card punch or a printer that is out of paper. The condition may also be

generated when the addressed device is in the not-ready state, in test mode, or not provided on the control unit.

2.7.2.3 Bus-Out Check. The bus-out check condition occurs when the I/O device or the control unit receives a data byte or a command byte with invalid parity over the I/O interface.

During writing, bus-out check indicates that a parity error was detected and that incorrect data may have been recorded at the I/O device. However, the condition does not cause the operation to be terminated prematurely, unless the operation is such that an error precludes meaningful continuation of the operation. No operation is initiated if the command code has a parity error.

2.7.2.4 Equipment Check. The equipment-check condition occurs when an equipment malfunction has been detected logically between the I/O interface and the I/O medium. On output operations, this malfunction may have caused invalid data to be recorded. Detection of equipment check stops data transmission and terminates the operation prematurely when the error prevents any meaningful continuation of the operation.

2.7.2.5 Data Check. The data-check condition occurs when invalid data has been detected by the control unit or I/O device. When signalled during execution of a write command, the data-check condition indicates that incorrect data may have been recorded by the I/O device. When signalled during execution of a read command, the data-check condition means that the control unit may have forced correct parity on the data transferred to the channel.

Data errors on reading and writing cause the operation to be terminated prematurely only when the errors prevent meaningful continuation of the operation (loss of synchronization).

2.7.2.6 Overrun. The overrun condition occurs when the channel fails to respond to the control unit in the anticipated time interval to a request for service from the I/O device. When the total activity initiated by the program exceeds the capability of the channel, an overrun may occur when data is transferred to or from a nonbuffered control unit or when the I/O device receives the new command too late during command chaining.

On an output operation, overrun indicates that data recorded at the device may be invalid. In these cases, data overrun normally stops data transfer, and the operation terminates as when 'stop' is indicated.

2.8 General System Considerations.

2.8.1 Interface Timeout Considerations. Except when the data-streaming feature is used, signalling over the I/O interface is specified to be direct-current-interlocked and is therefore not time-dependent. (See section 2.9.5 "Data-Streaming Feature.") Because the signalling may not always be time-dependent, a certain category of machine malfunctions may cause hangup of the channel and, unless detected and canceled, may result in hangup of the system. The timing considerations involved in determining malfunction cases are described in the following paragraphs.

All references in this section to particular time considerations represent the maximum permissible time in worst-case situations. All control units are designed for minimum signal-sequence-response times within the limitations of the circuit family used and the sequencing method required for particular I/O devices. The maximum time contributed by a control unit to an initial-selection sequence, or a short-busy sequence because of the control-unit circuitry, is no more than 32 microseconds. This restriction also applies to all byte-multiplex-mode connections. (See section 2.3.1 "Initial Selection Sequence.")

A particular interface signal sequence may take longer because of other factors, such as delays introduced by the channel or a delay because of the need for a burst-mode I/O device to capture the interface before reaching the specified record area of the recording medium (not greater than 500 milliseconds).

In addition to excessive delays that may occur within a particular signal sequence, excessive delays may exist between data-transfer sequences during burst-mode operation. The maximum delay between data-transfer sequences in burst mode is 500 milliseconds. The channel does not indicate a malfunction unless the delay between data-transfer sequences exceeds approximately 30 seconds. Activity on the I/O interface may be absent because of other conditions (reading a long gap on magnetic tape created by successive erase commands).

2.8.2 Propagation of Select Out. During an initial-selection sequence with an I/O-device address not assigned to a control unit or during a control-unit-initiated sequence, if the control unit does not require service or selection, 'select out' is propagated by the unit within 600 nanoseconds to meet performance requirements; in no case does it exceed 1.8 microseconds per control unit. (See "Initial-Selection Sequence" and "Control-Unit-Initiated Sequence" in this chapter.) If an initial-selection sequence with an address assigned to a control unit is attempted and (1) the addressed I/O device is unavailable and (2) 'select out' is to be propagated, the time required to propagate 'select out' may be the same as the time required for an initial-selection sequence for that control unit. This time, measured at the external cable connectors, extends from the rise of the incoming 'select out' to the rise of the outgoing 'select out.'

2.8.3 System Configuration.

2.8.3.1 Number of Units. In the general system configuration, as many as eight control units can be directly connected to a single channel I/O interface.

2.8.3.2 Internal Cabling. Except for 'select out' and 'select in,' the maximum allowable internal resistance (including all contact resistance) contributed by a channel or control unit for every signal line, is 1.0 ohm. The combined 'select out' and 'select in' maximum resistance for the control unit is 1.5 ohms.[6] The maximum resistance for 'select out' and 'select in' in a control unit usually occurs when the control-unit power is off and electrical bypassing is effective.

---

[6]For some older channels and control units designed before about 1970, the maximum allowable internal resistance (including all contact resistance) contributed by a channel or control unit for every signal line except for 'select out' and 'select in' is 2.0 ohms. The combined 'select out' and 'select in' maximum resistance for a control unit is 3.0 ohms.

For control units, the internal resistance is measured between the incoming and the outgoing pins on the external connectors. For channels, the measurement is made between the external connector pin and the corresponding channel driver or receiver.

With the exception of the 'select out'/'select in' line, the maximum signal delay measured between the external pins is 15 nanoseconds. The maximum skew between any two signal lines is 1 nanosecond. The delay for the bus lines does not exceed the delay for the tag lines.

The operational specification or vendor document for each control unit and channel must define, for each possible interface implementation, the value of its internal resistance for the signal lines and for the 'select out'/'select in' line pair.

2.8.3.3 External Cabling. The cable length available for the interconnection of channel and control units is primarily limited by the resistance to the interface lines which is contributed by the channel and the control units. For specific control units, the signal delays due to cable length require that the control unit be relatively close to the channel. However, the maximum allowable cable length is limited by total series resistance between drivers and receivers, if signal-cable delay requirements are met for each attached control unit.

The maximum external connector-to-connector cable length for unrestricted general systems configurations is determined by the combined internal resistance specification for 'select out'/'select in.' The maximum line resistance for the total 'select out'/'select in' loop is not greater than 52.5 ohms for worst-case conditions.

In certain customized installations where 'select out' is redriven at the end of the cable, the maximum cable length is determined by the 1.0-ohm internal control-unit and channel resistance specifications for the signal lines other than 'select out' and 'select in.' In this case, the maximum cable length can be calculated by using the 33-ohm driver-to-receiver maximum line resistance specification for worst-case conditions.

2.8.4 Offline/Online. Because in an offline mode it must not interfere in the operation between the channel and other control units on the same I/O interface, the control unit provides the following:

1. A logical bypass for 'select out.'

2. A gating off of all other line drivers to prevent interference on the I/O interface.

The control-unit transition to or from an off/online condition does not cause a machine malfunction. The following minimum conditions exist concurrently before the online-to-offline transition occurs:

1. The online/offline switch is set to OFFLINE.

2. The control unit is not actively communicating with the channel either for executing an I/O operation ('operational in' is active) or for signalling a busy condition ('status in' is active during a short-busy sequence).

3. The control unit is not currently disconnected from the l/O interface during the execution of an l/O operation involving that control unit.

4. No stacked, pending, or forthcoming status exists for this unit.

5. No command chaining is currently indicated for this unit.

Existence of the preceding conditions 2 through 5 ensures that no machine malfunctions occur when the operator throws the switch from ONLlNE to OFFLlNE.

2.8.5 Power Effects.

2.8.5.1 Steady State. The power-off state of any control unit does not affect any operations of other control units on the l/O interface. The control unit whose power is off provides an electrical bypass for 'select out,' and all of its interface driver and receiver circuits are prevented from interfering with the l/O-interface signals. The incoming 'select out' signal terminator at the receiver is disconnected when the signal is electrically bypassed. When power is off in all units, 'select out' is propagated back to the channel.

2.8.5.2 Transient and Spurious Signals. Each control unit is designed so that, if proper procedures are followed, the process of individually powering on or off does not cause the driver or receiver circuits of the control unit to generate noise on the I/O-interface signal lines.

2.8.5.3 Transient and Select-Out Pulse Splitting. Each control unit ensures that the process of electrically bypassing 'select out' before power changes does not interfere with 'select out' propagation, except for possible short discontinuities (less than 1.8 microseconds) because of contact bounce coincident with signal delay through parallel logic circuits. Also, each control unit ensures that 'select out' discontinuities, which may occur when another control unit on the I/O interface is being powered on or off, do not affect the propagation of 'select out.' This is accomplished by the use of a special latch circuit. The latch is turned on by the AND of 'select out' and 'hold out' and is reset by the fall of 'hold out.' The circuit is in series with the remaining selection logic in the control unit and causes a constant 'select out' signal within the control unit and, therefore, to all following control units, regardless of variations in the input 'select out' signal (figure 4).

2.8.5.4 Power Off/On Sequence Requirements. The combination of proper procedures and circuitry provides the following sequence of events for powering off:

1. Logical disconnection of the unit from the I/O interface. When the unit becomes logically disconnected, a green panel indicator (l/O interface disabled) in the proximity of the power-off control, comes on. (The control unit can become logically disconnected when it and all connected I/O devices have completed all I/O operations, when no status is pending or stacked, and when chaining is not indicated.) This ensures that no unfinished operations exist that can cause indication of machine malfunction when power is turned off. Note that when a control unit is logically disconnected from the l/O interface, all its drivers except 'select out' are logically gated off. A logical disconnect may be accomplished as

a result of going offline by use of the online/offline or the metering disable/enable switch.

2. Closing of the 'select out' bypass circuit (mechanical contact, K1, figure 4). The normal logical electronic bypass of 'select out' is still active when the control unit is logically disconnected.

3. Opening of the connection from the line to the 'select out' receiver terminator (mechanical contact, S1, figure 4).

4. Clamping of the interface driver gates to ground by means of a mechanical contact (S2, figure 4), if gated drivers are used to avoid transient signals on the I/O-interface lines.

5. Turning power off (remote/local power-control switch set to LOCAL). See FIPS 61 "Channel Level Power Control Interface."

For powering on, the sequence is reversed:

1. Turning power on (remote/local power-control switch set to LOCAL). During the power-on sequence, a power-on reset pulse automatically resets the control-unit circuitry, including resetting of the special 'select out' latch, regardless of 'hold out.'

2. Unclamping of driver gates. (Open mechanical contact, S2.)

3. Connection of the 'select out' receiver terminator (S1).

4. Opening of the 'select out' bypass circuit (K1).

5. Logical connection of the unit to the interface (panel indicator goes off).

If some method, such as automatic power sequencing rather than gated I/O-interface driver circuits, is used to eliminate transients on the signal lines, the steps in the sequence that refer to driver gates may be eliminated.

Note: The 'select out' bypass function (relay transfer) of a power-off or power-on sequence must be completed in one control unit attached to a channel before the 'select out' bypass function (relay transfer) is started in another control unit. Therefore, the 'select out' bypass function should be completely automatic or at least should be completely without interruption once it has started.

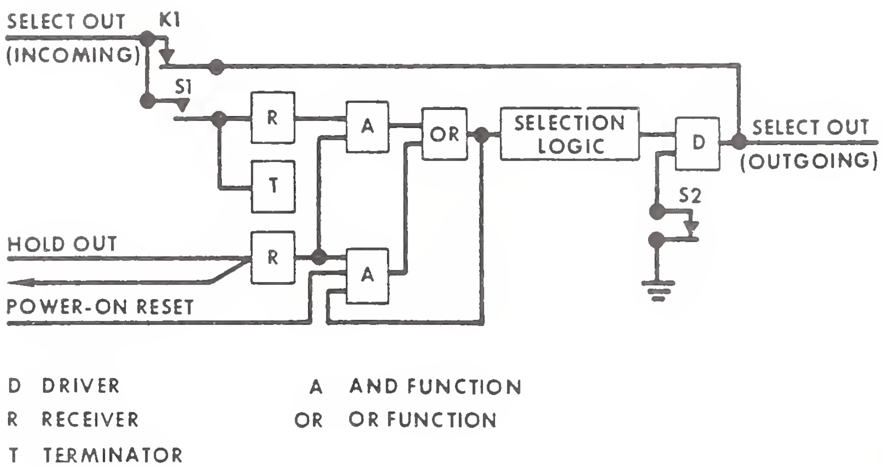

Figure 4. Representative Select Out/Hold Special Latch

2.9 Features. Unless otherwise noted, the definitions provided in section 1 and sections 2 through 2.8 apply when any of the features described here are used.

2.9.1 Bus-Extension Feature. This feature extends the interface by adding an extra bus to allow information transfer of two bytes in parallel instead of only one byte. The present bus is designated bus 0, while the additional bus is designated bus 1. If control units using different bus widths are attached to the same I/O interface, the control units using the wide bus must be attached first; the last control unit using the wide bus must terminate the bus not used by other control units.

The data lines of the additional bus are defined similarly to those of bus 0. In addition to data lines, 'mark in' and 'mark out' lines are defined for each bus cable, including the bus-0 cable, to indicate the presence of information on that bus. 'Mark in' and 'mark out' parity lines are also defined in the bus-1 cable to permit a validity check of the mark lines.

The number of mark lines active when 'address out,' 'address in,' 'command out,' or 'status in' are up indicates the number of bytes used for these functions. For output operations, the control unit must raise the 'mark in' lines before raising 'service in' or 'data in' to indicate the number of bytes required. The channel must then indicate by the 'mark out' lines the number of bytes provided on 'bus out' when 'service out' or 'data out' is raised. The bus width indicated during the early bus-width indication must be the maximum used by the control unit for that operation. (See "Early Data-Bus-Width Indication" in this chapter.) Transfers involving less than the maximum indicated bus width are permitted when the length of the block of information being transferred is not an integral multiple of the maximum transfer width. The partial transfer thus required is permitted as the last transfer of the block. All information transferred, including that transferred by a read-backward command, must be gated so that the first byte is on bus 0 and the second byte is on bus 1.

2.9.1.1 Mark-Out Lines. During sequences involving the transfer of information from the channel to the control unit, the 'mark out' lines are valid for the same period of time as the information on 'bus out' and indicate the number of bytes provided by the channel on 'bus out.'

For sequences involving the transfer of information from the control unit to the channel, the 'mark out' lines indicate the number of bytes that the channel is accepting. The 'mark out' lines are valid from the rise of the outbound tag line until the resulting fall of the inbound tag line.

A channel capable of extended-bus operation ensures that the number of 'mark out' lines up, including 'mark out parity,' is odd when the I/O device with 'address out' up is addressed. Additionally, if the control unit with which the channel is communicating is operating with extended-bus capabilities, the channel maintains odd mark-line parity whenever 'service out,' 'data out,' or 'command out' is up.

2.9.1.2 Mark-In Lines. During sequences involving the transfer of information from the control unit to the channel, the 'mark in' lines are valid for the same period of time as the information on 'bus in' and indicate the number of bytes provided by the control unit on 'bus in.'

For sequences involving the transfer of information from the channel to the control unit, the 'mark in' lines indicate the number of bytes required by the control unit for that sequence and are valid from the rise of the inbound tag line until the rise of the corresponding outbound tag line.

The number of 'mark in' lines up, including 'mark in parity,' must be odd whenever a control unit is operating with extended-bus capabilities and 'service in,' 'data in,' 'address in,' or 'status in' is up. Control units operating with the one-byte interface do not necessarily activate any 'mark in' lines, except for command-retry indication. (See section 2.9.3 "Command-Retry Feature.")

2.9.1.3 Early Data-Bus-Width Indication. The control unit indicates to the channel the bus width to be used for an operation by raising the 'mark in' lines during the initial-selection sequence. These 'mark in' lines are valid from the rise of 'operational in' to the fall of 'address out.'

2.9.2 I/O-Error-Alert Feature. When a malfunction that affects the continued-execution capability of a control unit occurs, 'disconnect in' may be raised to alert the channel of the condition. An example of such a condition is one in which a microcoded control unit is communicating with the channel at the time a microcode storage error is detected. Such a control unit may be unable to complete an interface sequence properly. Another example is one in which the control unit, while currently disconnected from the I/O interface, recognizes a microcode storage error while communicating with an attached I/O device. 'Disconnect in' can be raised by a control unit only when it is connected to the channel (that is, it has 'operational in' up). If the control unit is currently disconnected from the I/O interface, a control-unit-initiated sequence is performed to establish a connection before 'disconnect in' is raised. When 'disconnect in' is used during a control-unit-initiated sequence, 'disconnect in' rises at the control unit no earlier than 250 nanoseconds after 'address in' rises, or subsequent to the rise of 'command out' at the control unit.

The channel, in response to 'disconnect in,' performs a selective reset. 'Disconnect in' does not fall before the reset nor remain up longer than 100 nanoseconds after the fall of 'operational in.'

2.9.3 Command-Retry Feature. Command retry is a channel and control-unit procedure that can cause a command to be retried by the channel. The command-retry procedure is initiated by the control unit with a unique combination of status bits and the use of the 'mark 0 in' line.

A control unit may request the retry of a command in order to recover from a transient error or when conditions existing at either the control unit or I/O device prevented execution of the command when it was previously issued.

A channel, upon accepting a request for command retry, repeats the execution of the channel program, beginning at the last command executed.

2.9.3.1 Command-Retry Sequence. If, during execution of a command, the control unit encounters a condition requiring retry, the control unit requests command retry by raising 'mark 0 in' and 'status in' while presenting unit check and status modifier together with (1) channel end (meaning the control unit or the device is not yet ready to retry the command), or (2) channel end and device end (meaning the control unit and device are prepared for immediate retry of the command). Device end, if not presented with channel end, is presented later, when the control unit is ready to retry the command.

The channel acknowledges the request for command retry by indicating command chaining. If device end accompanies the command-retry request, the channel immmediately initiates an initial-selection sequence and reissues the previous command. If device end does not accompany the command-retry request, command chaining is indicated, but the retry is not immediately performed. When device end or device end with status modifier is presented to the channel, command chaining is indicated, and an initial-selection sequence is performed to reissue (1) for device end, the previous command and (2) for device end with status modifier, a new command from the modified sequence of commands. (See 2.6.4 "Status Modifier.")

A channel indicates a refusal to perform a command retry by accepting the status byte without indicating chaining or by stacking the status byte and thus not indicating command chaining. The stacked byte is treated as any stacked status. When the stacked status is subsequently accepted by the channel, command chaining is again not indicated, and the command-retry procedure is not performed.

2.9.4 High-Speed-Transfer Feature. Some control units have data-transfer-rate requirements that exceed the capabilities provided through the use of only 'service in' and 'service out.' When longer I/O-interface-cable lengths are used, the high-speed-transfer feature permits data transfers to take place at higher data rates than those achievable when 'service in' and 'service out' are used. Specifically, higher data rates are made possible by alternating 'data in' with 'service in' and 'data out' with 'service out,' respectively. However, the use of this feature does not require this alternation of tag lines. Situations may occur for example, where successive alternations of only 'data in' and 'data out' are performed.

The high-speed-transfer feature may also be used to allow placement of a control unit at a greater distance from the channel than is otherwise possible.

2.9.4.1 Data In.

Note: When the data-streaming feature is used, this "Data In" section does not apply. (See section 2.9.5 "Data-Streaming Feature.")

'Data in' is a tag line to the channel from all control units with this feature. It is used to signal the channel that the selected control unit requires the transmission of a byte of information. The nature of the information depends on the I/O operation and the I/O device. The channel responds to the rise of 'data in' by raising either 'data out' or 'command out.'

During execution of an I/O operation specified by read, read-backward, and sense commands, the control unit places a data byte on 'bus in' and raises 'data in.' During execution of a write command or control command, the control unit raises 'data in,' and the channel places a data byte on 'bus out.' When 'data in' is alternated with 'service in,' 'data in' may rise when 'service out' is raised in response to 'service in.' However, 'data in' is not considered valid until 'service in' is dropped. Similarly, 'service in' may rise when 'data out' is raised in response to 'data in'; however, 'service in' is not considered valid until 'data in' is dropped.

The conditions that apply to 'service in' concerning overrun also apply to 'data in.'

2.9.4.2 Data Out.

Note: When the data-streaming feature is used, this "Data Out" section does not apply. (See section 2.9.5 "Data-Streaming Feature.")

'Data out' is a tag line from the channel to all attached control units and is used in response to the rise of 'data in.' During read, read-backward, and sense operations, the rise of 'data out' indicates that the channel has accepted the information on 'bus in.' During execution of a write or control command, the rise of 'data out' indicates that the channel has placed on 'bus out' the data requested by 'data in.'

When 'data out' is sent in response to 'data in' during execution of an I/O operation specified by a read, read-backward, or sense command, 'data out' rises after the channel accepts the information on 'bus in.' In these cases, the rise of 'data out' indicates that the information is no longer required to be valid on 'bus in' and is not associated with any information on 'bus out.' When 'data out' is sent in response to 'data in' during execution of a write or control command, the rise of 'data out' indicates that the channel has provided the requested information on 'bus out.' 'Data out' remains up until the fall of 'data in.'

2.9.5 Data-Streaming Feature. The data-streaming feature allows channels and control units to operate at a high rate through longer cables than would be possible with the high-speed transfer feature alone (see section 2.9.4 "High-Speed Transfer Feature.") Channels which implement the data-streaming feature must also implement the high-speed transfer feature. Channels which implement the data-streaming feature shall be capable of supporting control units, which also implement this feature, and which have a data transfer rate of up to 3 megabytes per second, when connected through up to 122 meters (400 feet) of interface cable, without the use of the bus-extension feature (see section 2.9.1 "Bus Extension Feature").

Whether a control unit can use this feature is determined by the data-rate and cable-length requirements it must meet. The normal direct-current-interlock method described elsewhere in this standard imposes the less critical timing considerations on a control unit. The data-streaming feature requires more precise timing implementations but offers higher data rates that are independent of the cable lengths.

The data-streaming feature is applicable only for read and write commands; it is not used for sense or control commands. For the duration of each read or write command, 'service in' is alternated with 'data in.'

The data-streaming feature requires that both control units and channels have no greater than ±22.5 nanoseconds of internal skew between the rise of the service or data tag and the rise of any associated bus-bit signal.

All timings specified in the definition for out-tags or 'bus out' are measured at the channel tailgate. All timings specified in the definition for in-tags or 'bus in' are measured at the control-unit tailgate.

Since control units which do not implement the data-streaming feature may not observe the convention defined in the immediately following paragraphs to distinguish that the transfer is to use data-streaming (i.e., they may initiate a dc-interlocked transfer by raising 'data in'), I/O channels which implement the data-streaming feature shall be capable of supporting control units and devices which do not implement the data-streaming feature. It is permissible to partition the I/O channel device address space into data-streaming and non-data-streaming segments, or to use switch settings to distinguish data-streaming and non-data-streaming devices. If setup or restrictions on the use of the address space are required, then this shall be defined in the vendor's documentation.

Those control units that can transfer data by using the data-streaming feature signal the channel which method of data transfer is desired each time a data transfer is initiated. The following convention is used to identify the desired method of data transfer.

1. If the control unit initiates data transfer by raising 'data in,' data is transferred using the data-streaming feature.

2. If the control unit initiates data transfer by raising 'service in,' data is transferred using the normal dc-interlock method.

The channel must know if the control unit has the data-streaming feature so as to properly interpret the in tags.

The maximum frequency at which data is transferred by using the data-streaming feature is defined in the vendor's documentation for the I/O device and for the channel implementing this feature. The maximum frequency is expressed as the inverse of the minimum time from the rise of 'service in' or 'data in' until the next rise of 'service in' or 'data in,' respectively. Note that each cycle transfers two bytes. In addition, the channel must ensure that the rate of out-tag responses take place within $\pm15\%$ of the rate that the in-tag requests are received. That is, the interval from the rise of 'service out' or 'data out' until the next rise of 'service out' or 'data out,' respectively, must be within $\pm15\%$ of the interval from the rise of 'service ' or 'data in' until the next rise of 'service in' or 'data in,' respectively, that was actually received by the channel.

Note: When the data-streaming feature is used during data transfer, data must be valid on the appropriate bus for the specified amount of time. If the receiver of the data is delayed in sampling the bus, data may be lost. Thus, channels and I/O devices that previously did not have a time-dependency on the receipt of data (such as buffered I/O devices) might now overrun when this feature is implemented.

2.9.5.1 Data Transfer While Data Streaming. When data transfer is performed by using data streaming, the rise and fall of 'service in' or 'data in' is independent of the rise and fall of 'service out' or 'data out,' respectively. Furthermore, for 'service in' or 'data in,' the uptime or downtime is not less than 270 nanoseconds, respectively. For 'service out' or 'data out,' the uptime or downtime for either out tag is not less than 180 nanoseconds.

During execution of an I/O operation specified by a read or read-backward command, data is valid on 'bus in' from 22.5 nanoseconds after the rise of 'service in' or 'data in' until 247.5 nanoseconds or greater after the rise of 'service in' or 'data in.' Also, the overlap between 'service in' and the next consecutive 'data in' or between 'data in' and the next consecutive 'service in' is not greater than 22.5 nanoseconds.

During execution of an I/O operation specified by a read or read-backward command, 'service out' or 'data out' rises in response to 'service in' or 'data in,' respectively indicating that the channel has accepted the information on 'bus in.'

When, during execution of an I/O operation specified by a write command, 'service out' or 'data out' rises in response to 'service in' or 'data in,' respectively, the

channel has placed data on 'bus out' and, after the appropriate 100-nanosecond delay, has raised the corresponding out tag for not less than 180 nanoseconds. Data is valid on 'bus out' from 100 nanoseconds before the rise of 'service out' or 'data out' until 100 nanoseconds or more after the rise of 'service out' or 'data out,' respectively.

For I/O operations specified by read- or write-type commands, once an in tag is recognized and acted upon by the channel, the out-tag signal is dropped only after the corresponding in-tag signal has fallen. 'Service out' and 'data out' may be up concurrently and are not up with any other out-tag except when 'command out' is used to signal stop or during an interface-disconnect sequence control, when 'address out' may be up.

When the channel does not respond in time to a preceding request for service, an error condition is recognized, and the I/O operation is then ended by the control unit.

Notes:

1. If a channel determines that the I/O device is requesting service faster than the channel can respond, the channel does not respond, thus forcing the I/O device to detect an error condition. When the I/O device detects the error condition, it initiates the ending sequence. When the ending sequence is initiated by the control unit, the channel responds by accepting or stacking the status as in a normal ending sequence.
2. Not requiring a normal dc-interlock for the first byte of data transfer can cause problems for channels with buffered I/O devices attached. Nonbuffered I/O devices have some mechanical delay before the first data-transfer sequence, and current channels rely upon this time to initialize their data-transfer hardware. However, a buffered I/O device may require data transfer immediately and possibly cause an overrun condition to be recognized.

   The channel may delay the initiation of data transfer in one of two ways. The channel may (1) delay data transfer by the use of 'suppress out' or (2) delay dropping of the out-tag immediately following the completion of either the initial-selection sequence or the control-unit-initiated sequence, as applicable.

   When the channel delays the initiation of data transfer by the use of 'suppress out,' the channel does so as defined in "Suppress Data While Data Streaming."

   When the channel delays the initiation of data transfer by delaying the dropping of the out-tag immediately following the completion of either the initial-selection sequence or the control-unit-initiated sequence, it follows this procedure:

   a. After an initial-selection sequence has been completed, the channel delays the fall of 'service out' until it is ready to transfer data. Thus, an I/O device may not assume that the channel is capable of data transfer until it has recognized the fall of 'service out' in response to the rise of 'status in.' If 'suppress out' is up when 'service out' falls, indicating suppress data, an I/O device that initiates data transfer may create an overrun condition.

b. During a control-unit-initiated sequence, the channel delays the fall of 'command out' until the channel is ready to transfer data. Thus, an I/O device may not assume that the channel is capable of performing a data-transfer sequence until it has recognized the fall of 'command out' in response to the rise of 'address in.'

2.9.5.2 Stop/Command Out While Data Streaming. The stop sequence control is indicated by 'command out' in response to 'service in' or 'data in.'

'Command out' in response to 'service in' or 'data in' is used to signal the I/O device that the channel is ending the current operation. On receipt of the stop signal, the I/O device proceeds to its normal ending point without raising either 'service in' or 'data in' The I/O device remains busy until the ending status is accepted by the channel.

When data streaming is used and data transfer is to be stopped, 'command out' is raised in response to the first 'service in' or 'data in' that is raised after the channel determines to end the operation. 'Command out' remains up until the fall of the associated 'service in' or 'data in.' However, the interval between the rise and fall of 'command out' is not less than 180 nanoseconds.

Following the signalling of stop, one or more additional 'service in' or 'data in' signals may, because of propagation delay, be received by the channel and are responded to by the appropriate 'service out' or 'data out.' For each such signal, during read or read-backward operations, the channel does not check 'bus in' for correct parity. For each additional 'service in' or 'data in' received during write operations, the channel responds with the appropriate 'service out' or 'data out' signal, and 'bus out' has a byte of all zeros with undefined parity. 'Bus out' is not checked for parity or decoded by the control unit after stop has been signalled. The total number of channel responses equals the number of I/O-device requests. If they do not compare equal, the control unit sets the unit-check status indicator and provides sense data specifying data check.

Note: Because 'command out' is substituted for 'service out' or 'data out' when stop is signalled, 'command out,' 'service out,' and 'data out' may be detected in any paired combination and be active concurrently.

2.9.5.3 Suppress Data While Data Streaming. For the data-streaming control unit, the definition for 'suppress data' is identical to that described elsewhere in this manual, except that if the control unit disconnects from the channel during data transfer, it recognizes 'suppress out' for the first data byte after reselection. A control unit that provides the data-streaming feature does not initiate data transfer until 'suppress out' drops or until further delay would result in the detection of an overrun condition, whichever is sooner.

When the data-streaming feature is used, the channel may, because of propagation delay, recognize the rise and fall of one or more 'service in' or 'data in' signals after 'suppress out' has been raised. Thus, for read or read-backward operations, the channel is capable of accepting the additional bytes of data after 'suppress out' has been signalled.

Note: Channel designs must take into account that, when suppress data is signalled, additional data-transfer requests (depending upon the I/O-interface cable length) may be received after 'suppress out' is raised. Therefore, 'suppress out' should be raised prior to the time when data transfer should stop in order to ensure that all data requested is handled at the required rate.

2.9.5.4 Interface Disconnect While Data Streaming. When the data-streaming feature is used, interface disconnect is signalled in the same manner as when the direct-current-interlock method is used. Also, when data streaming is used, the rules for maintaining the validity of data during data transfer on 'bus in' for an I/O operation specified by either a read or read-backward command or on 'bus out' for an I/O operation specified by a write command are unchanged during the signalling of interface disconnect. However, when data is being transferred by using the data-streaming feature, one or more additional 'service in' or 'data in' signals may still be received by the channel after the signalling of interface disconnect. The channel does not respond to these signals with 'service out' or 'data out.' For each 'service in' or 'data in' received during read or read-backward operations, the channel does not check 'bus in' for correct parity. For each additional 'service in' or 'data in' received during write operations, the channel does not respond with the corresponding 'service out' or 'data,' respectively. When 'operational in' drops, the channel may drop 'address out' to complete the interface-disconnect sequence control.

2.9.5.5 Selective Reset While Data Streaming. Following the signalling of 'selective reset,' one or more additional 'service in' or 'data in' signals may, because of propagation delay, be received by the channel. In this situation, the additional 'service in' or 'data in' signals received are not responded to by the rise of the appropriate 'service out' or 'data out.' For each additional 'service in' or 'data in' during execution of an I/O operation specified by either a read or read-backward command, the channel does not check 'bus in' for correct parity. For each additional 'service in' or 'data in' received during execution of an I/O operation specified by either a read or read-backward command, the channel does not check 'bus in' for correct parity. For each additional 'service in' or 'data in' received during execution of an I/O operation specified by a write command, the channel does not respond with the corresponding 'service out' or 'data out,' respectively.

2.9.5.6 Response-Time Requirements While Data Streaming. When data streaming is used, during data transfer the control unit may or may not receive the coresponding response to the last 'service in' or 'data in' in the expected time interval. To avoid potential hangup conditions, the control unit may recognize an error condition and present unit-check status if at least 8 microseconds has elapsed from the rise of the in-tag that is awaiting an out-tag response.

CHAPTER 3.

Electrical Specifications

3.1 Multiple Drivers and Receivers. Up to ten receivers must be able to be driven by one driver. The driver must be located at one of the extreme ends. Up to ten drivers must be able to be 'dot' ORed to drive one receiver. The receiver must be located at one of the extreme ends. There will be no more than one driver and/or receiver per line per control unit. This limitation also applies to the channel.

Note: An end-of-line driver or receiver may be placed beyond the terminator. In this case, the distance between the end-of-line driver or receiver and the terminator must be less than six inches. Receivers must be spaced at least three feet apart. No minimum requirement is set for the spacing between drivers. No minimum requirement is set for the spacing between a terminator and driver or receiver if the terminator is placed on the outermost end of the line. The maximum stub length from the line to a driver or receiver on the circuit card is six inches.

3.2 General Electrical Requirements.

3.2.1 Voltage Levels. There are two logical levels. A dc line voltage of +2.25 volts or more denotes a logical one state, and a dc voltage of +0.15 volt or less denotes a logical zero state. These voltages are relative to the driver ground.

3.2.2 Cable. All lines must have a characteristic impedance of 92 $\pm$ 10 ohms and, with the exception of 'select out,' must be terminated at each extreme end in their characteristic impedance by a terminating network. (For 'select out'/'select in,' see section 3.5 "Electrical Specifications for Select Out Circuitry.") Cable length may be limited by special conditions but is never to exceed a maximum line resistance of 33 ohms. The 33 ohm line resistance includes all contact resistance, internal cable resistance, and interunit cable resistance.

3.2.3 Terminating Networks. The terminating network must present an impedance of 95 ohms $\pm$ 2.5 percent between the signal line and ground, and must be capable of dissipating 390 milliwatts.

3.2.4 Ground Shift and Noise. The maximum noise (measured at the receiver input) coupled onto any signal line must not exceed 400 millivolts. The maximum allowed ground shift, between an active driver and any receiver of the same interface line, is 150 millivolts. Therefore, the maximum shift (coupled noise plus ground shift) allowed on any line is 550 millivolts.

The logical levels defined in section 3.2.1 "Voltage Levels" under section 3.2 "General Electrical Requirements" and the receiver threshold levels specified in section 3.3.1 "Receivers" under section 3.3 "Interface Circuit Requirements" allow for this 550 millivolt shift. That is, a negative noise pulse of 400 millivolts coupled with a positive receiver ground shift of 150 millivolts occurring during a one state (2.25 volts minimum) guarantees a receiver input of 1.7 volts or more. (See diagram A.)

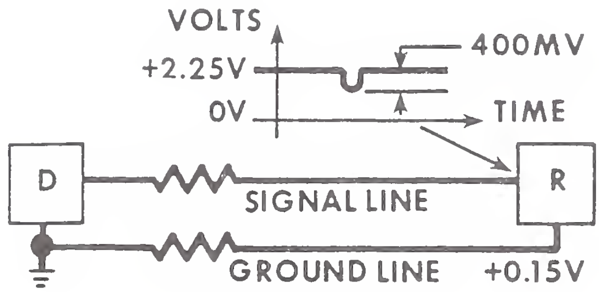

DIAGRAM A. Negative Noise

Also, a positive, noise pulse of 400 millivolts coupled with a negative receiver ground shift of 150 millivolts occurring during a zero state (0.15 volt maximum) guarantees a receiver input of 0.7 volt or less. (See diagram B.)

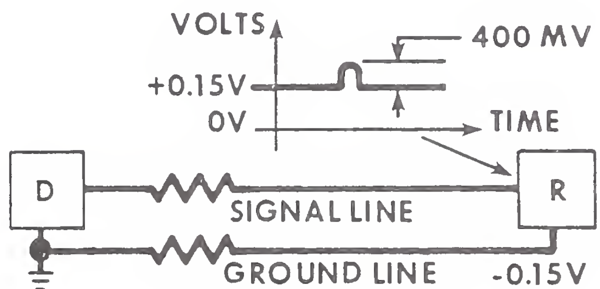

DIAGRAM B. Positive Noise

Note: The noise measurements are made at the input to the receiver. A combination of the dc level and ac noise must not exceed 0.7 volt for the down level and must not be less than 1.7 volts for the up level.

3.3 Interface Circuit Requirements.

3.3.1 Receivers. An input voltage (relative to receiver circuit ground) of 1.7 volts or more is interpreted as a logical one; an input of 0.70 volt or less is interpreted as a logical zero. An open-circuited input is interpreted as a logical zero.

The receiver should not be damaged by:

1. A dc input of 7.0 volts with power on in the receiver.

2. A dc input of 6.0 volts with power off in the receiver.

3. A dc input of -0.15 volt with power on or off.

To reduce the loading effect on the line, the receiver input resistance is larger than 7,400 ohms across the input voltage range of +0.15 volts to +3.9 volts, and the negative receiver input current does not exceed -0.24 milliampere at an input voltage of +0.24 milliampere at an input voltage of +0.15 volt. (See diagram C for the definition of receiver-current polarity.)

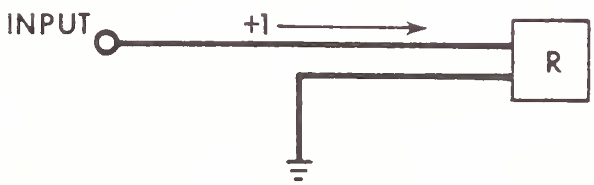

DIAGRAM C. Receiver

3.3.2 Drivers. In the logical zero state:

1. The output voltage must not exceed 0.15 volt at a load of +240 microamperes. (See diagram D for conventional current polarity definition.)

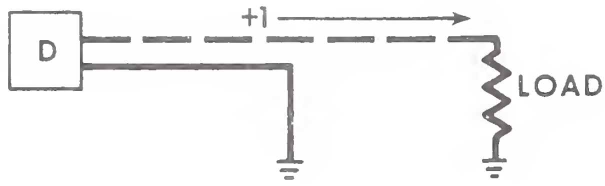

DIAGRAM D. Current Polarity

In the logical one state:

1. The ouptut voltage must be 3.11 volts or more at a load of +59.3 milliamperes (two terminators, ten receivers).

2. The output voltage must not exceed 5.85 volts at a load of +30 microamperes (one receiver, no terminator).

3. The output voltage must not exceed 7.0 volts at a load of +123.0 milliamperes during an overvoltage internal to the drivers.

Drivers must be designed to ensure that no spurious noise is generated on the line during a normal power-up or power-down sequence. For the driver, this may be accomplished by one of the following methods.

1. Sequencing the power supplies.

2. Building noise suppression into the circuit.

3. Providing an externally controlled gate. (See diagram E.)

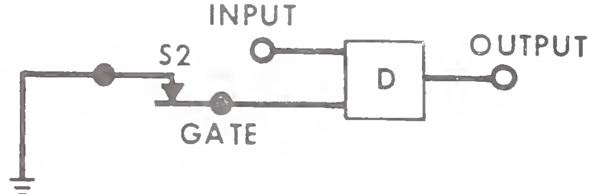

DIAGRAM E. Driver Gate

For normal power-down sequence:

1. Logically ensure that the driver is in the zero state.

2. Close contact S2. (See diagram E.)

3. Turn power off.

For a normal power-up sequence:

1. Ensure that contact S is closed.

2. Turn on power.

3. Logically ensure that the input level will cause the driver output to be in the zero state.

4. Open contact S.

3.3.3 Fault Conditions. A grounded signal line must not damage drivers, receivers, or terminators. With one driver transmitting a logical one, loss of power in any single circuit driver, receiver, or terminator on the line must not cause damage to other components. With both terminators connected, line operations must not be affected by power off in any drivers or receivers on the line.

3.4 Typical Circuits. Figures 5 through 7 show representative circuits used to drive, receive, and terminate the lines between the channel and attached control units.

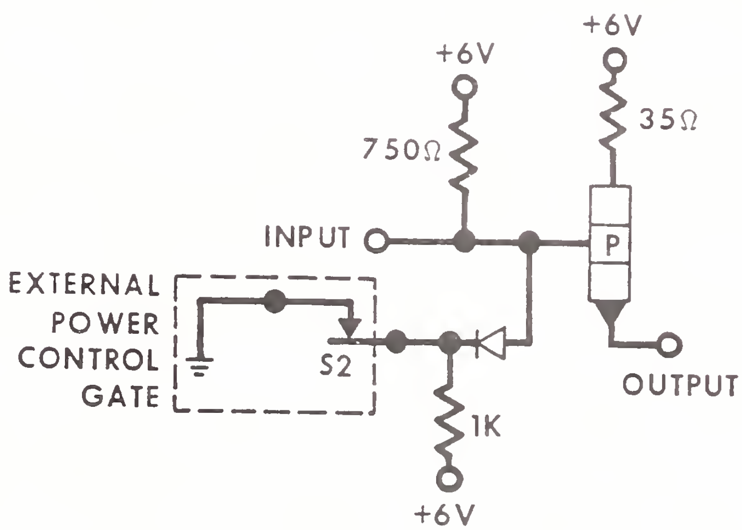

FIGURE 5. Line Drivers

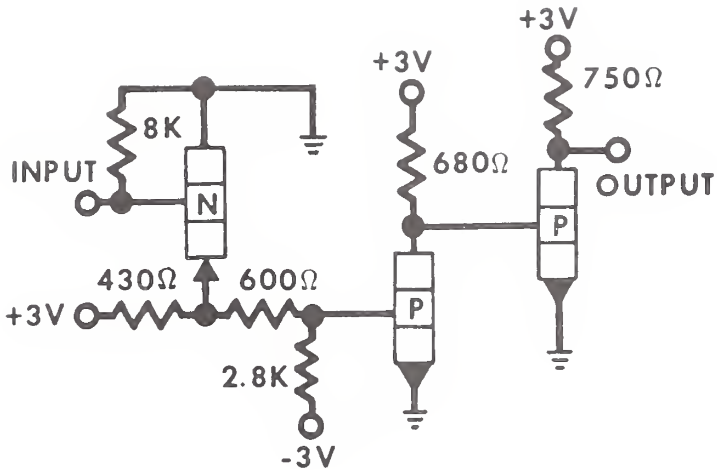

FIGURE 6. Line Receivers

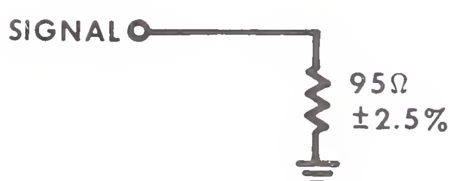

FIGURE 7. Line Terminator

3.5 Electrical Specifications for Select Out Circuitry.

3.5.1 General. The 'select out' line has a single-driver to single-receiver configuration, with only the receiver end of the line terminated in the characteristic impedance. A dc line voltage of 1.85 volts or more denotes a logical one state, and a dc line voltage of 0.15 volt or less denotes a logical zero state. These voltages are relative to the driver ground.

Note: Because of the nature of the 'select out'/'select in' line, negative noise tolerance has been neglected.

All electrical requirements specified in section 3.2 "General Electrical Requirements" that are not redefined in this section are also applicable to 'select out.'

3.5.2 Receiver. The 'select out' receiver must satisfy all requirements given in section 3.3.1 "Receivers."

3.5.3 Driver. The 'select out' driver must be capable of withstanding a short-circuit-to-ground output condition, while in either the logical one or zero state, without damage to the driver circuit. For the logical zero state:

1. The output voltage of a 'select out' driver must not exceed 0.15 volt at a load of 1 milliampere.

For a logical one state:

1. Output voltage of a channel driver or the driver of a control unit contained within a channel frame must exceed 3.9 volts at a load of 41 milliamperes.

2. The output voltage of a control unit driver not contained within a channel frame must exceed 3.7 volts at a load of 41 milliamperes.

The output voltage of a 'select out' driver should not exceed:

1. 5.8 volts at a load of 0.3 milliampere.

2. 7.0 volts at a load of 72 milliamperes during an overvoltage internal to the driver.

3.5.4 Terminator. A 95 ohm $\pm$ 2.5 percent, 390 milliwatt terminator to ground must be placed at each receiver for each line segment along the 'select out'/'select in' path, including the receiver end of 'select in' located in the channel.

The driver end of each segment of 'select out'/'select in' must not be terminated, including the driver end of 'select out' located in the channel. Also the jumpered 'select out' or 'select in' and the bypassed 'select out' or 'select in' path is not terminated.

3.5.5 Measuring I/O-Interface Timings. The definition of the I/O interface specifies several different timings. These timings are measured in the following manner.

3.5.5.1 Uptime. The uptime measured from when the signal goes above the least-positive-up level (LPUL) to when it goes below the LPUL. (See figure 8.) The LPUL for an active driver at its output is +3.11 volts at +59.3 mA. When the signal is measured at the unit's external connector, the voltage drop caused by the +59.3 mA flowing through the internal cable should be taken into account and the measurement voltage of +3.11 adjusted accordingly. This LPUL value ensures that the input to all receivers on the interface meets or exceeds the requirements of the timing diagrams when measured at the receiver's LPUL of +1.70 volts.

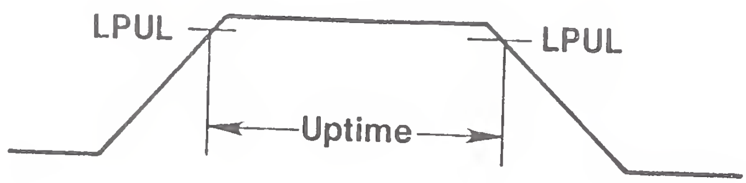

Figure 8. Uptime

3.5.5.2 Downtime. The downtime is measured from the point at which the signal goes below the most-positive-down level (MPDL) to the point at which it goes above the MPDL. (See figure 9.) The MPDL for an active driver is +0.15 volt. This MPDL value ensures that the input to all receivers on the interface meets or exceeds the requirements of the timing diagrams when measured at the receiver's MPDL of +0.70 volt.

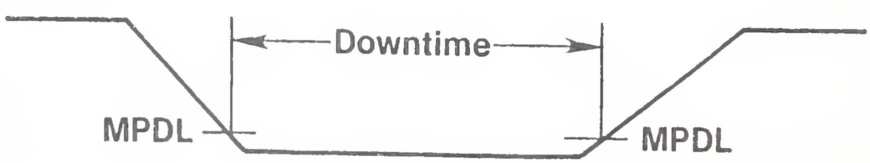

Figure 9. Downtime

3.5.5.3 Valid Data. Valid data is measured from when the signal goes above the LPUL for a logical one or below the MPDL for a logical zero.

3.5.5.4 Overlap Not Greater Than. The point of measurement of an "overlap not greater than" is from when the first signal goes below the LPUL to when the second signal goes above the LPUL. (See figure 10.)

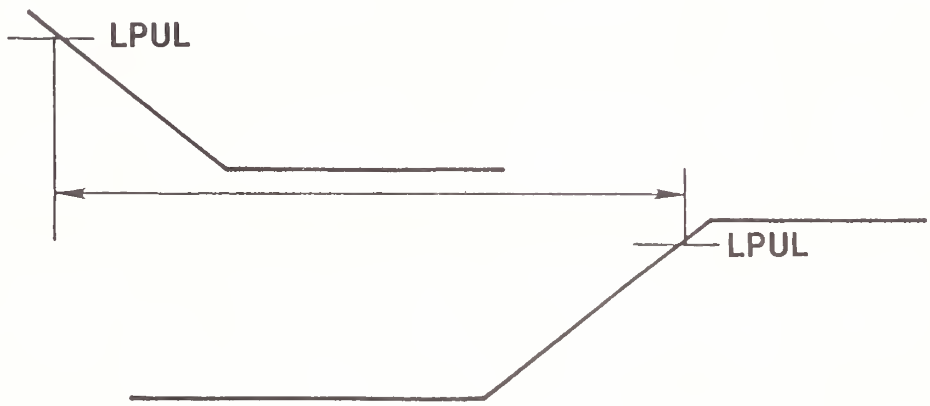

Figure 10. Overlap Not Greater Than

3.5.5.5 100-Nanosecond Delay. The 100-nanosecond delay (see sections 2.2.1.1 "Bus Out" and 2.2.1.2 "Bus In") is measured from when the data is valid until the corresponding out tag goes above the MPDL. (See figure 11.)

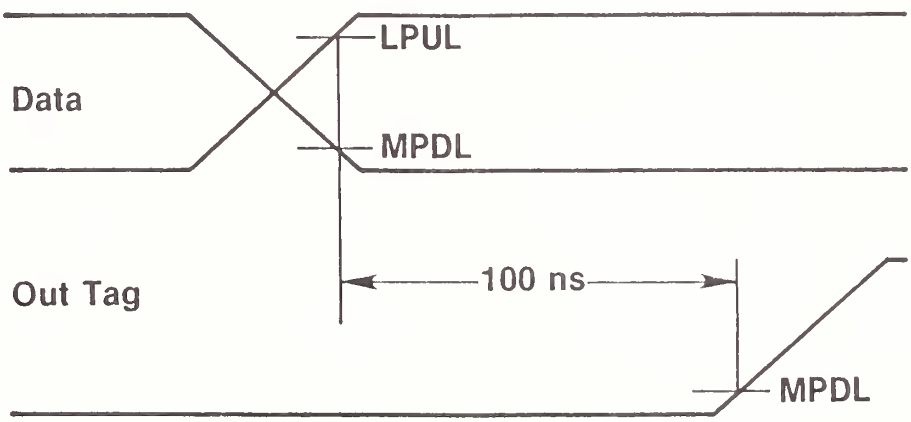

Figure 11. 100-Nanosecond Delay

3.5.6 Interface-Connector Pin Assignments. Pin assignments for the I/O interface are shown in figure 12 as viewed from the connector side of the channel and control-unit tailgates. Two interface cables (three cables for the optional two-byte interface) connect the channel to the first of a group of control units. Signals are assigned corresponding pin numbers on the channel and control-unit connectors (see Chapter 4 "Mechanical Specifications").

CHAPTER 4.

Mechanical Specifications

4.1 Cabling.

4.1.1 Cable Halves. A screwdriver is required to join the cable halves. The connecting screw is spring-loaded to prevent damage to mating block contacts. Blocks should be aligned to ensure proper parallel contact mating. After the blocks are keyed, push the screw forward to engage the insert; then tighten securely.

4.1.2 Mounting. In normal applications, the blocks are mounted in the horizontal plane to provide a smoother bend into the coaxial cables (see figure 13). Vertical mounting produces unusual bending configurations, requiring careful routing and strain relieving of the external cables.

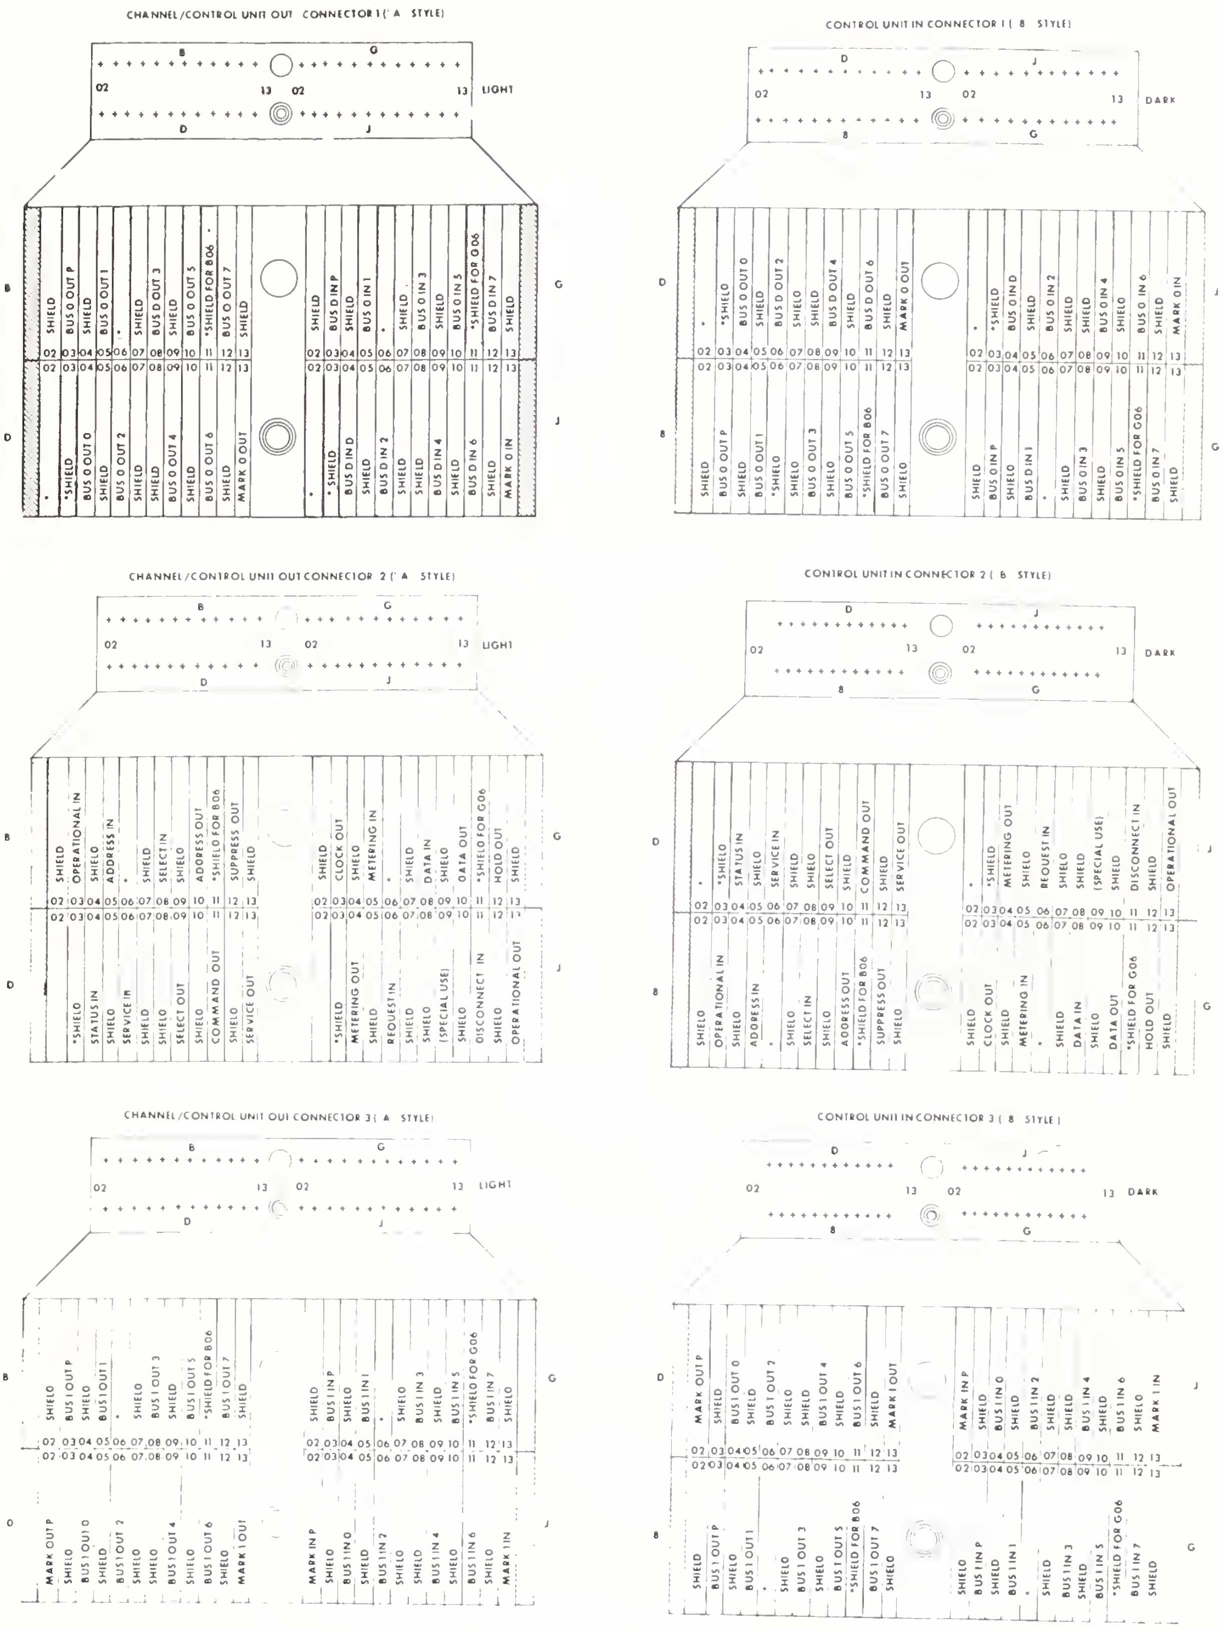

*Reserved for future use

Note: Shading indicates those signals passed only with 48-position cable

FIGURE 12. I/O Interface Connector Pin Assignment

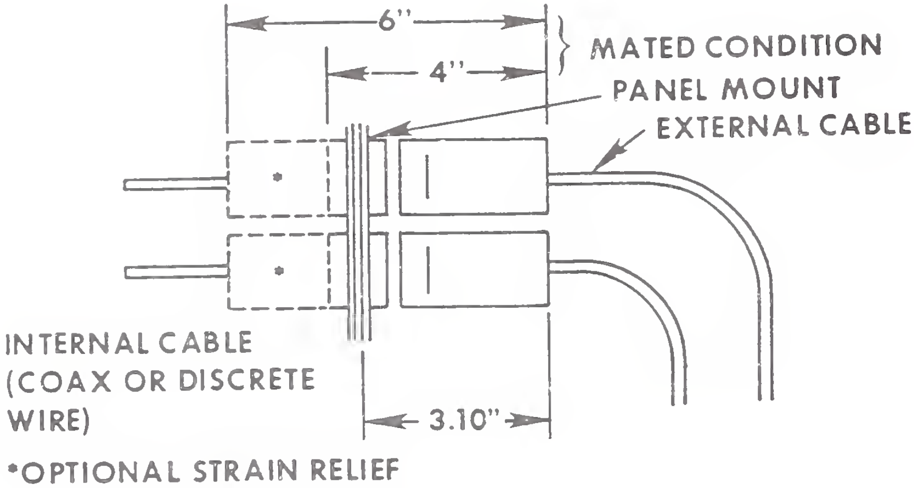

FIGURE 13. Typical Mounted Cable Halves

Figure 14 shows an exploded view of the panel mounting.

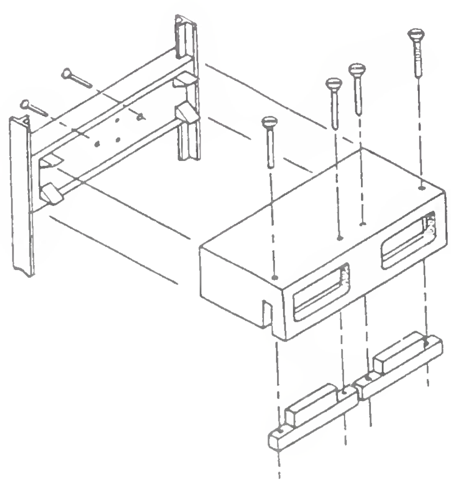

FIGURE 14. Typical Panel Mount with Optional Strain Relief

4.1.3 Offline Utilization. This cable connector has the unique feature of being able to mate cable-half to cable-half for offline utilization or for the physical bypassing of machine units (see figure 15).

When mating cable halves, it is only necessary to use the screw on the "B-" style connector. This allows the "A-" style connector's screw to fall within the empty

insert location in the "B-" style block and makes for easier connection. Typical connections are shown in figures 15 and 16.

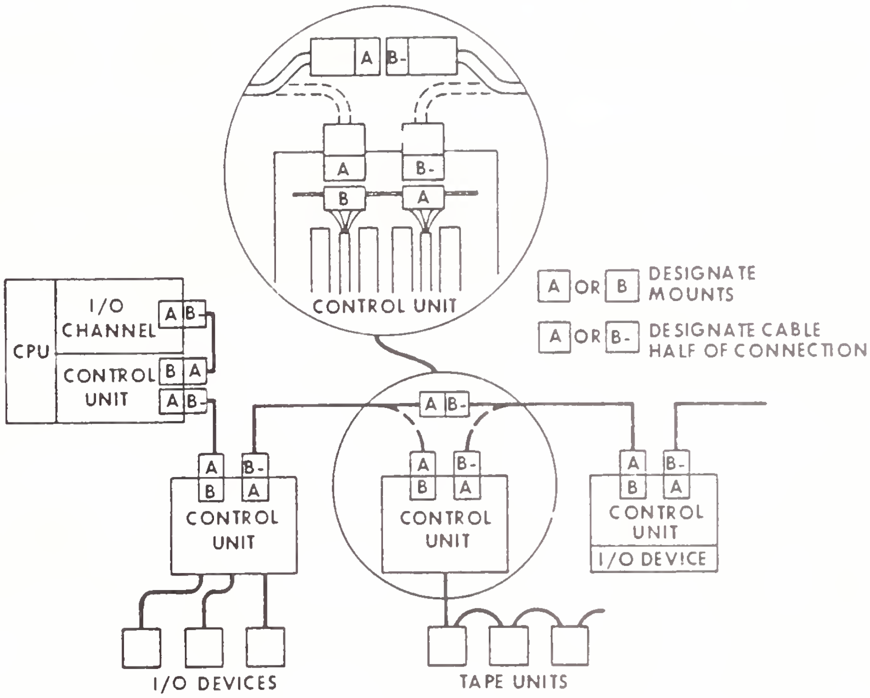

FIGURE 15. Offline Utilization

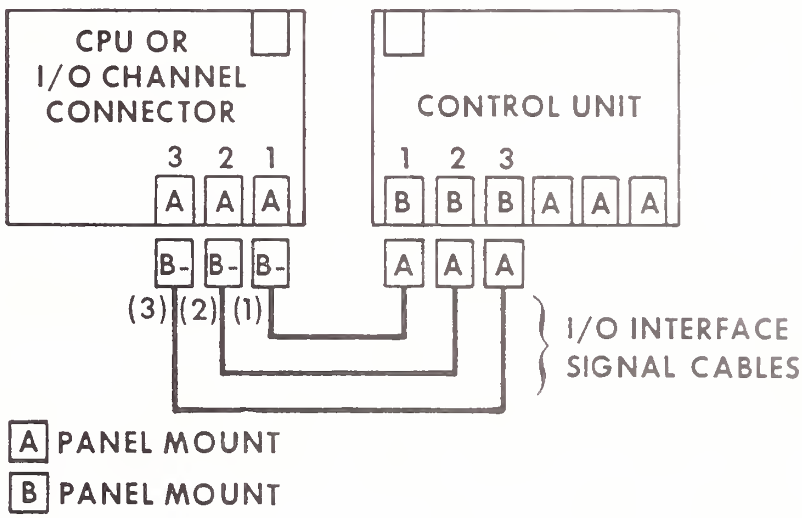

FIGURE 16. Typical Connection

4.2 Connectors.

4.2.1 Connector Block. Three styles of connector blocks are available: "A," "B," and "B-." The "A" and "B" designations are used to identify proper mating arrangements since the physical hardware is identical. The two styles are differentiated by the color coding of the blocks: the "A" style is a light color and the "B" style is a dark color (see figure 17). The "B-" style block is the same as the "B" style but does not have a threaded insert.

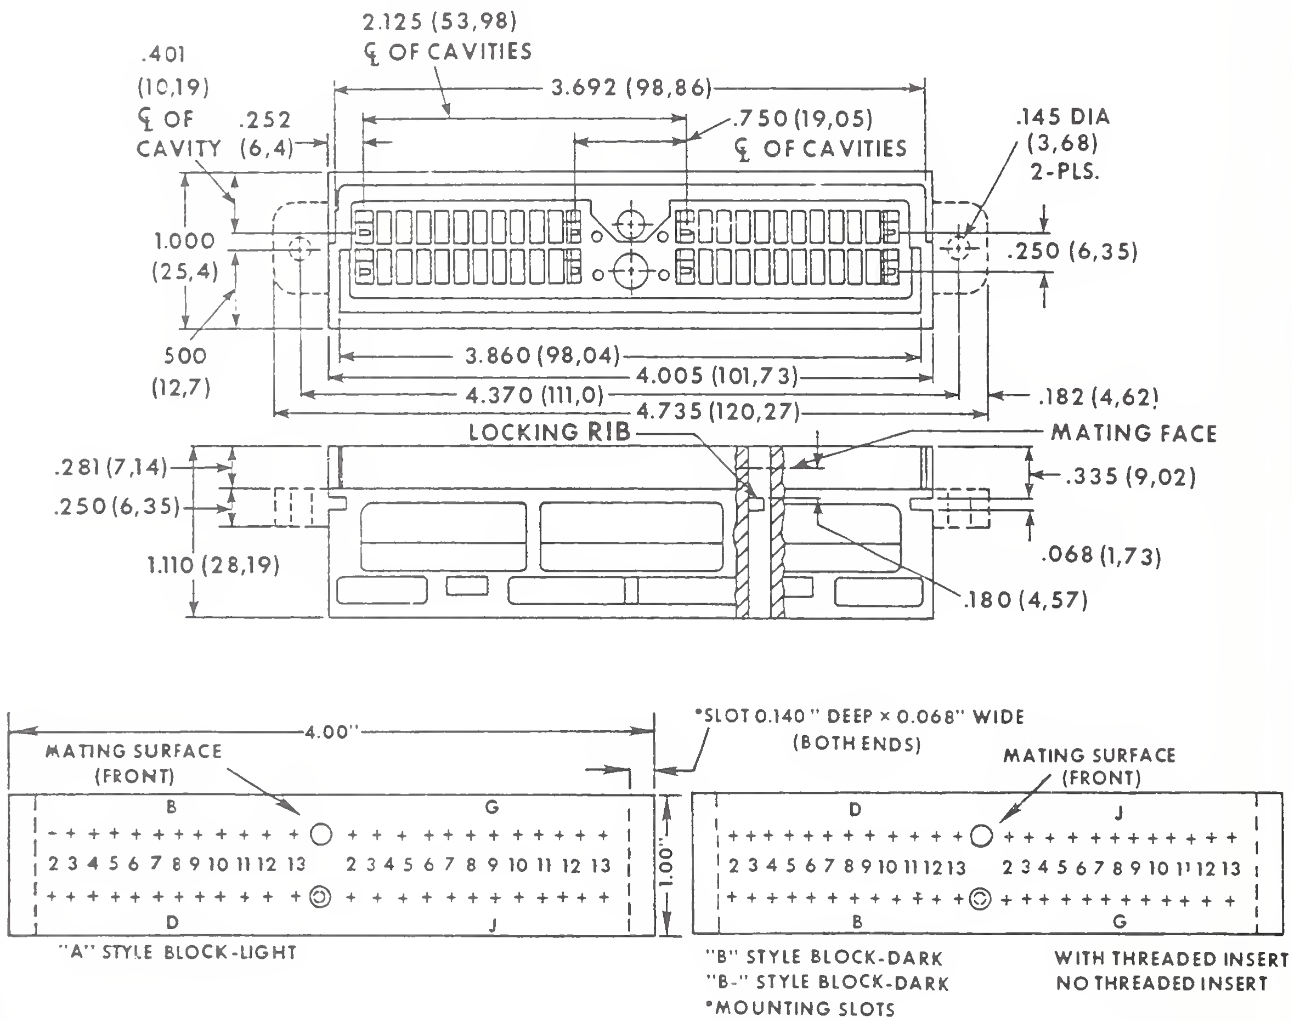

FIGURE 17. Connector Blocks and Contact Location

When mating connectors, care should be taken to prevent accidental mismating of two "A" or two "B" style connectors since letter positions would then be transposed. Connectors of the same color must never be mated. Both the "A" and the "B" style blocks are used for panel mounts. The "A" style may also be used for a cable end. The "B-" style is used only as a cable end (refer to section 4.1.3 "Offline Utilization").

4.2.2 Capacity. Forty-eight (48) individual positions (contacts) are provided for in the connector blocks. Application of the connector is limited only by the number of coaxial wires (shielded wires or twisted pairs) used and the method used to ground the shields of the coaxial wires.

4.2.3 Contacts. The contact is a hermaphroditic, gold-plated phosphor bronze, dual-mating surface contact or equivalent (see figure 18).

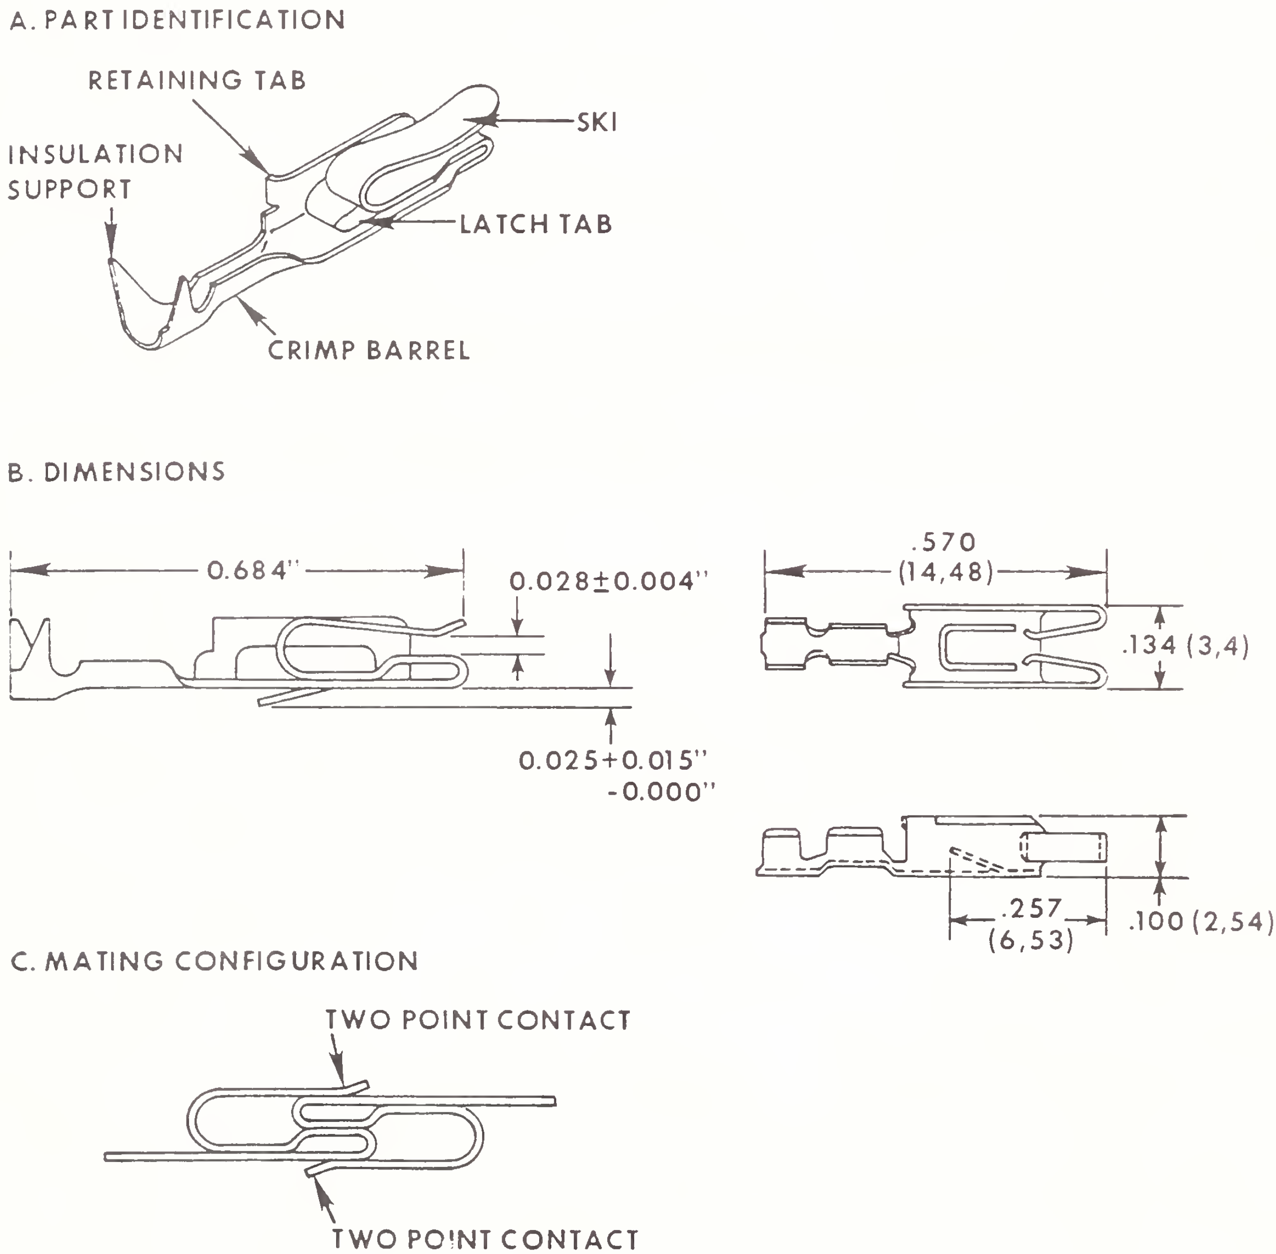

FIGURE 18. Typical Contacts

4.2.3.1 Voltage Ratings. The maximum voltage rating of this connector is 24 volts ac or dc.

4.2.3.2 Current Rating. The maximum continuous current rating of each contact is 6 amperes. The contacts are not intended for interrupting current.

4.2.3.3 Resistance. The termination-to-termination resistance (includes 2 crimps and mated contacts) will not exceed:

1. 0.020 ohms when installed on #22 AWG and larger wire.

2. 0.030 ohms when installed on #24-#26 AWG wire.

3. 0.040 ohms when installed on #28-#32 AWG wire.

4.2.3.4 Insulation Resistance. The contact-to-contact insulation resistance is 100 megohms (minimum) measured at a test potential of 100 volts dc, after exposure of one hour at a temperature of 38 degrees Centigrade and 85-90 percent relative humidity.

4.2.3.5 Grounding. All surfaces of the connectors are nonconductive plastic; therefore, no grounding is necessary.

4.3 Terminators. The I/O interface line termination is provided by the assemblies shown in figure 19. There is one tag and one bus assembly required, the difference being the handling (turn-around) of the 'select out' line.

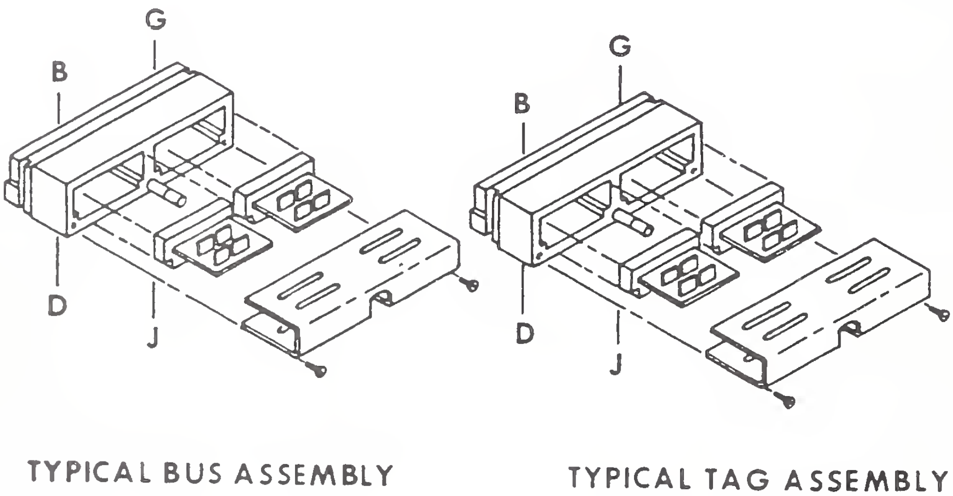

FIGURE 19. Terminator Assemblies

APPENDIX A.

FLOW DIAGRAMS

(This Appendix is not a part of the Standard, Specifications for I/O Channel Interface but is included for information purposes only.)

Note: The terms "Byte Multiplexer Channel", "Block Multiplexer Channel", and "Selector Channel" are used to describe possible implementations of this standard.

AT DECISION BLOCKS INVOLVING SIGNAL LINES, THE QUESTION TO BE RESOLVED IS --- IS THE LINE LOGICALLY UP.

THE PRESENCE OF CH (CHANNEL) OR CU (CONTROL UNIT) IN A BLOCK, SPECIFIES RESPONSIBILITY FOR THE ACTION TAKEN OR DECISION MADE.

THE START-POLL DECISION IS USUALLY MADE ON THE BASIS OF REQUEST IN BEING UP.

REQUEST IN IS RAISED WHEN A CONTROL UNIT HAS A NEED FOR SERVICING AND IS DROPPED WHEN THE NEED IS SERVICED.

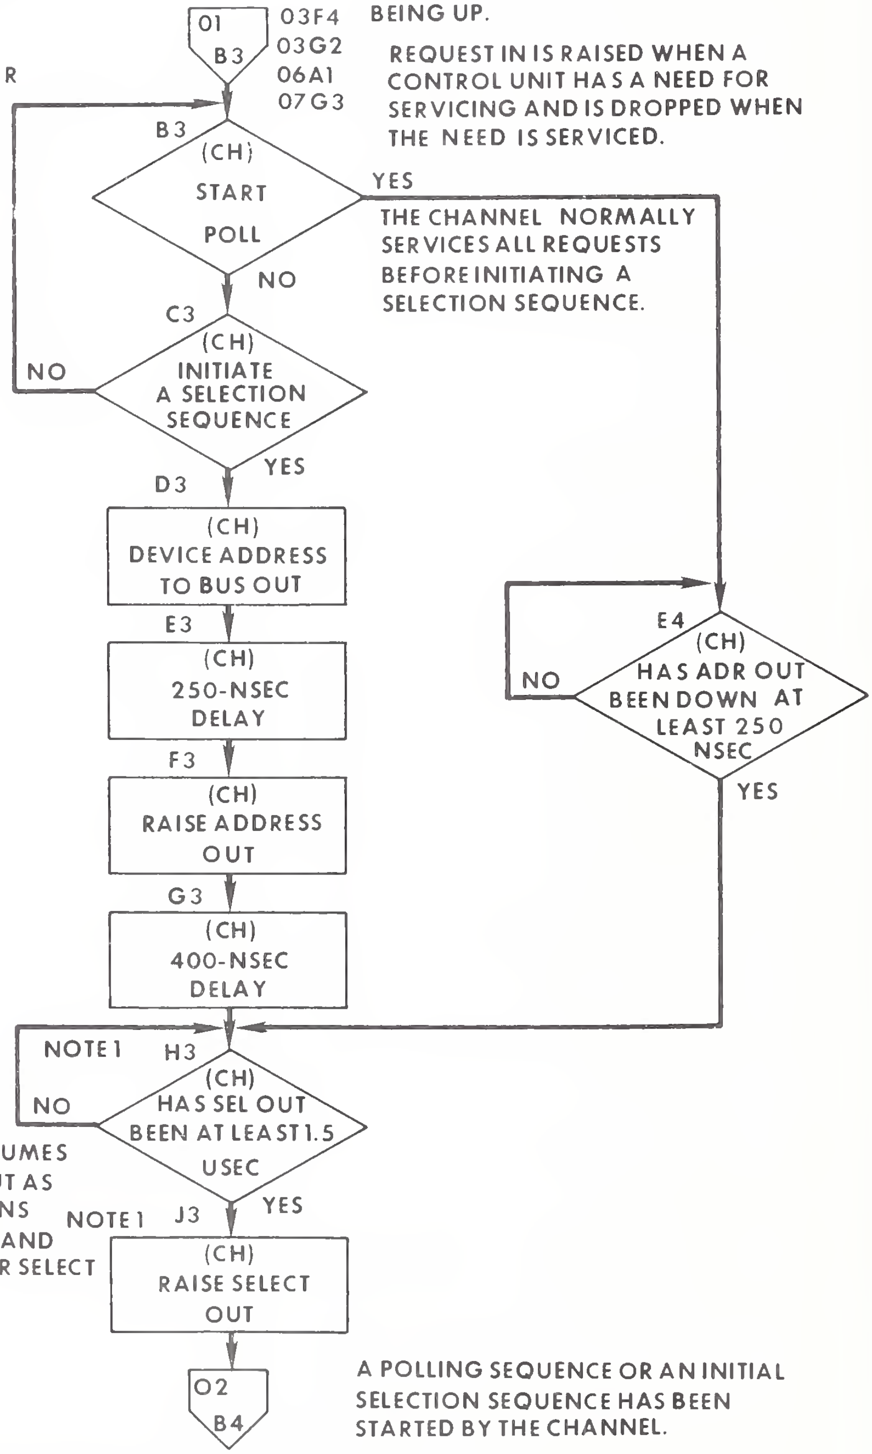

NOTE 1

ALL MENTION OF SELECT OUT ASSUMES PROPER OPERATION OF HOLD OUT AS WELL. THUS, SELECT OUT UP MEANS SELECT OUT AND HOLD OUT UP, AND SELECT OUT DOWN MEANS EITHER SELECT OUT OR HOLD OUT DOWN.

A POLLING SEQUENCE OR AN INITIAL SELECTION SEQUENCE HAS BEEN STARTED BY THE CHANNEL.

FLOW DIAGRAM 1. *Initiation of Polling or Selection*

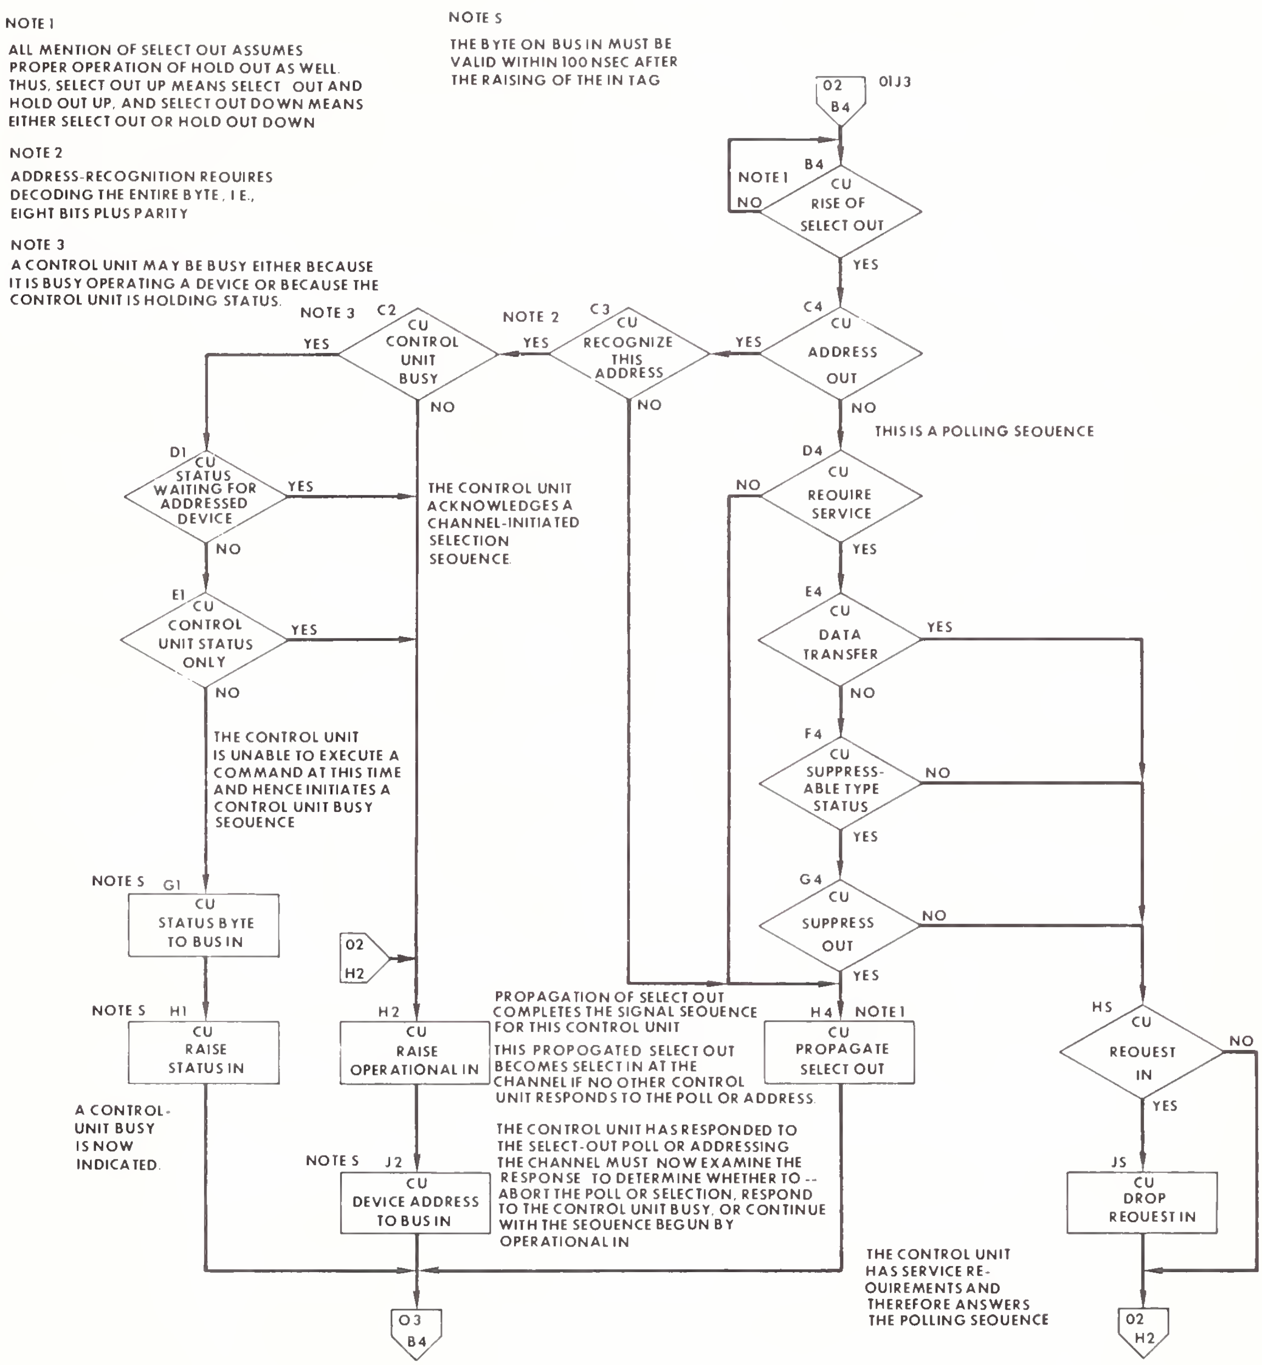

FLOW DIAGRAM 2. *Control Unit Response to Select Out (Part 1 of 2)*

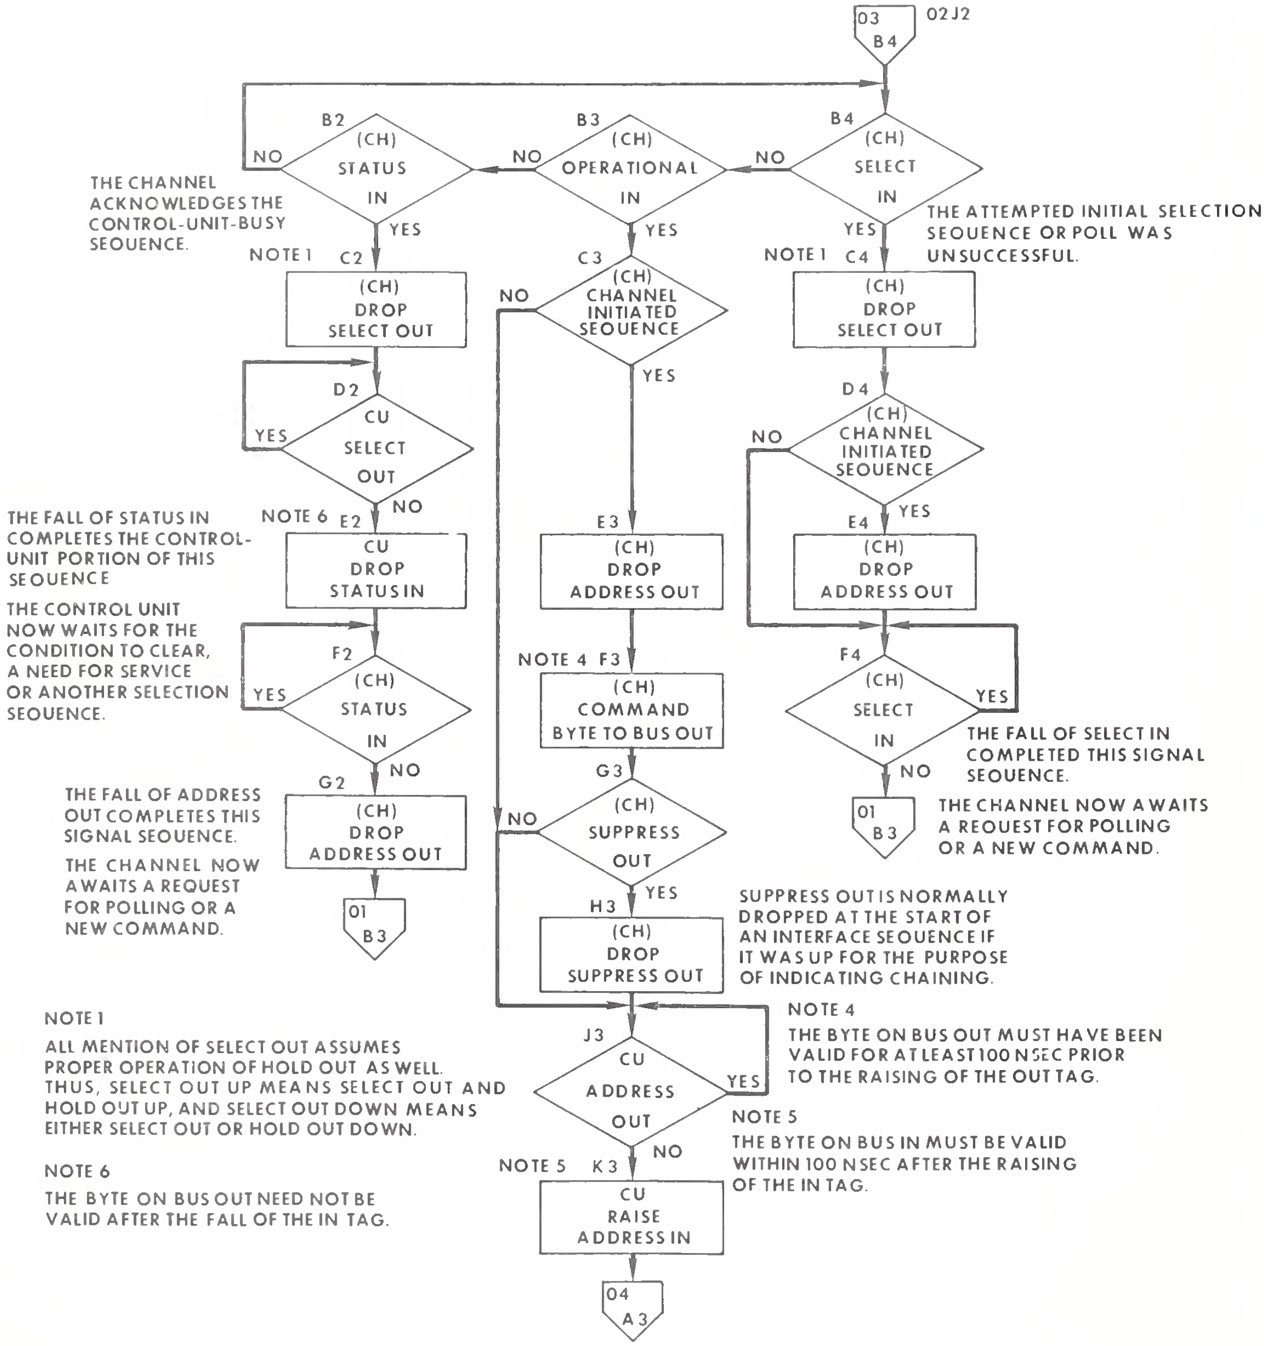

Flow Diagram 2. *Control Unit Response to Select Out (Part 2 of 2)*

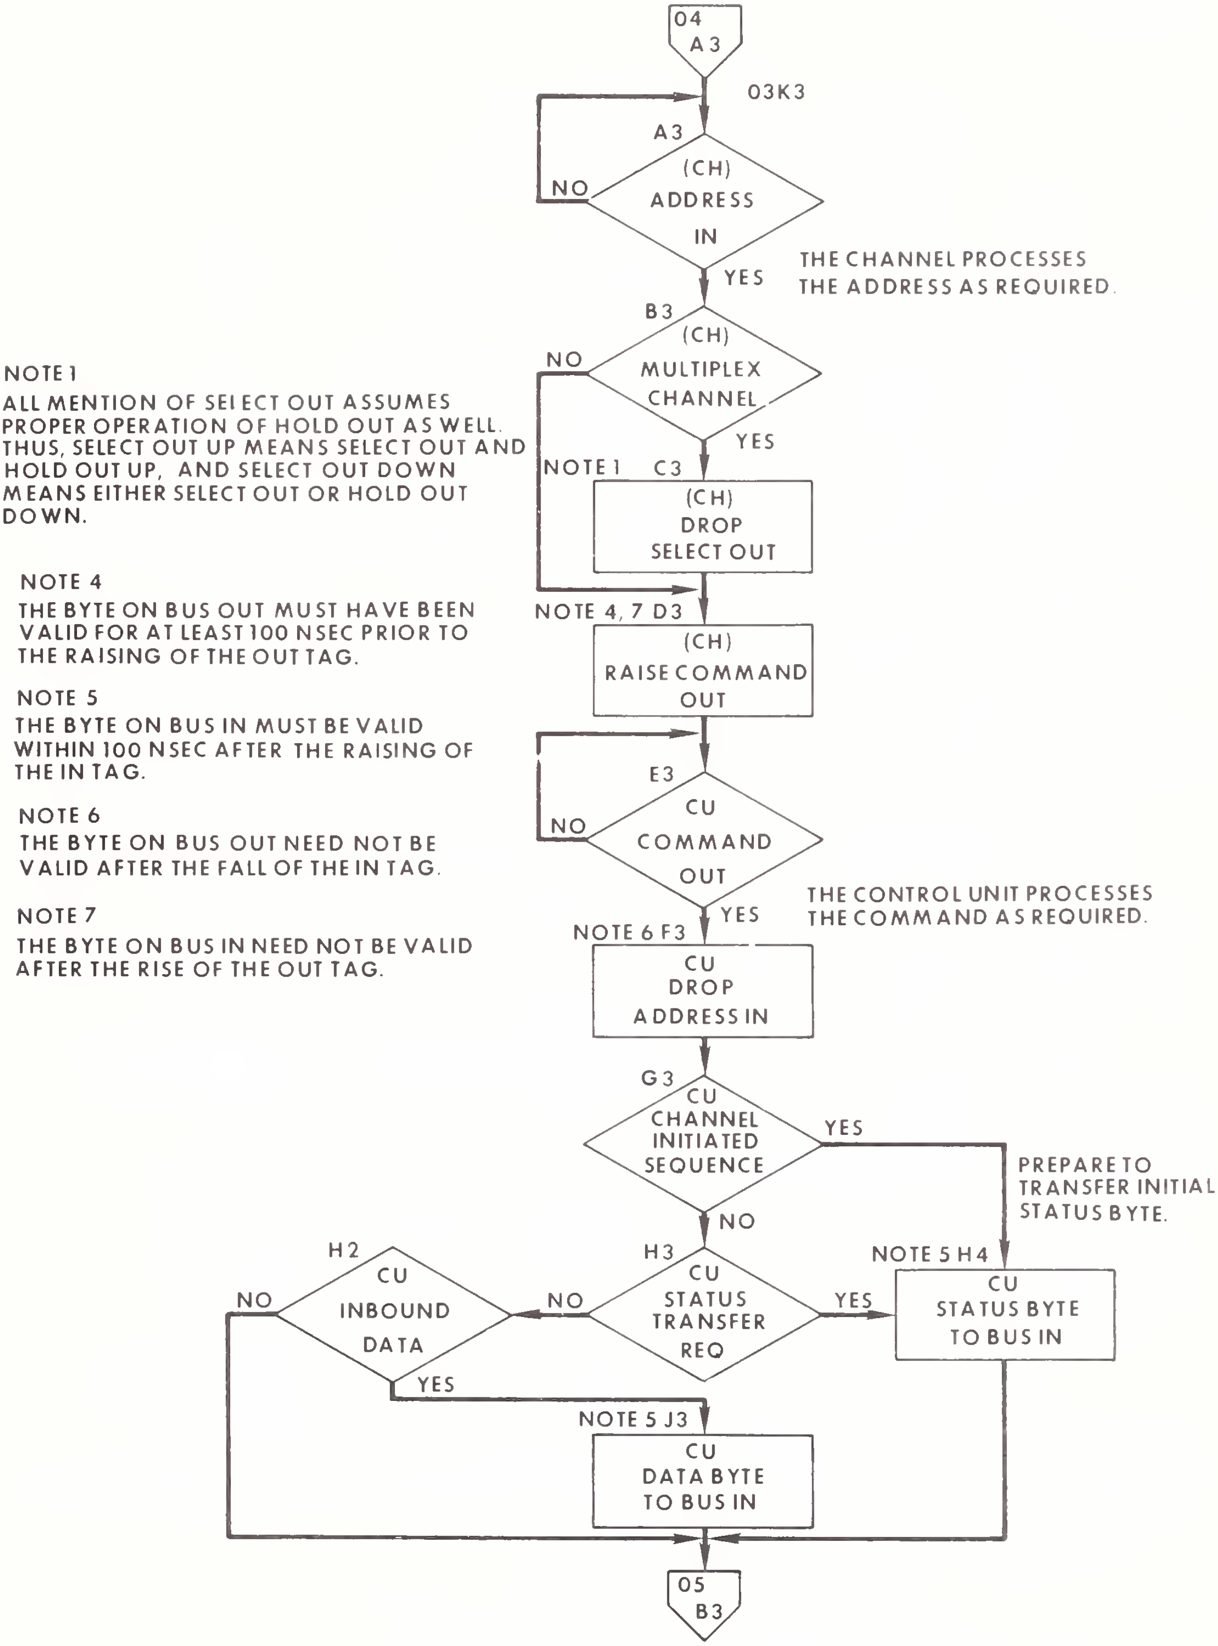

FLOW DIAGRAM 3. *Command Transfer*

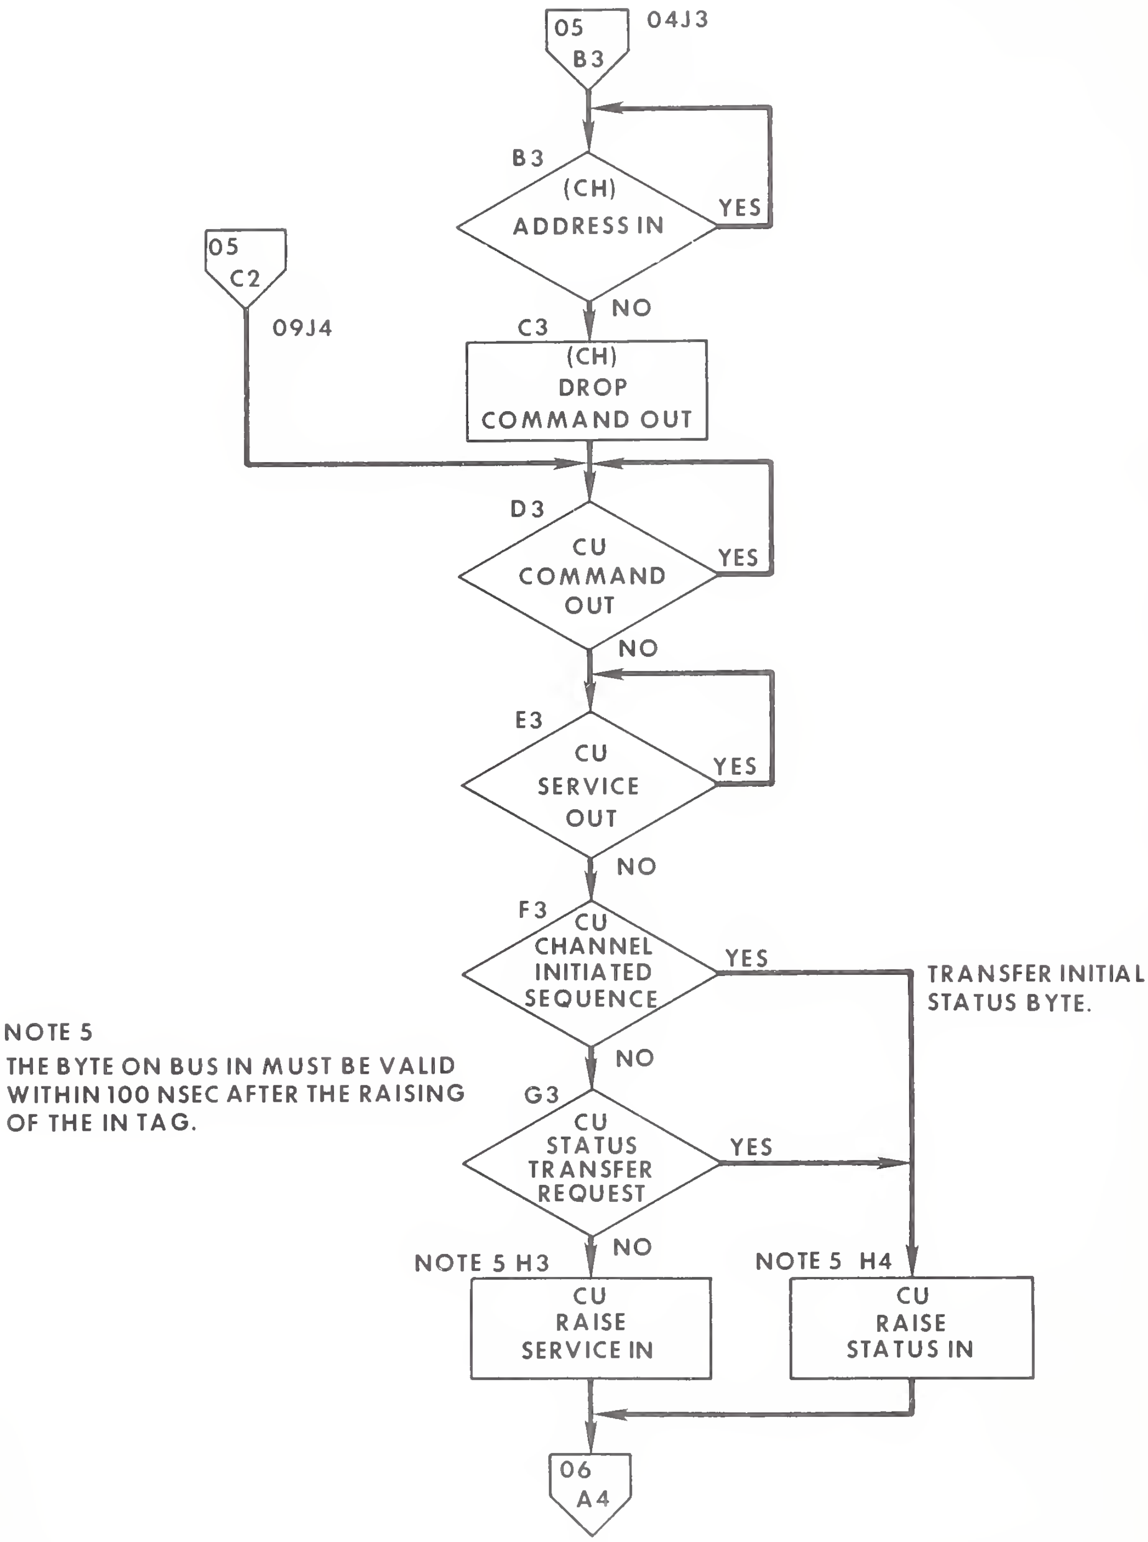

FLOW DIAGRAM 4. *Status/Data Presentation*

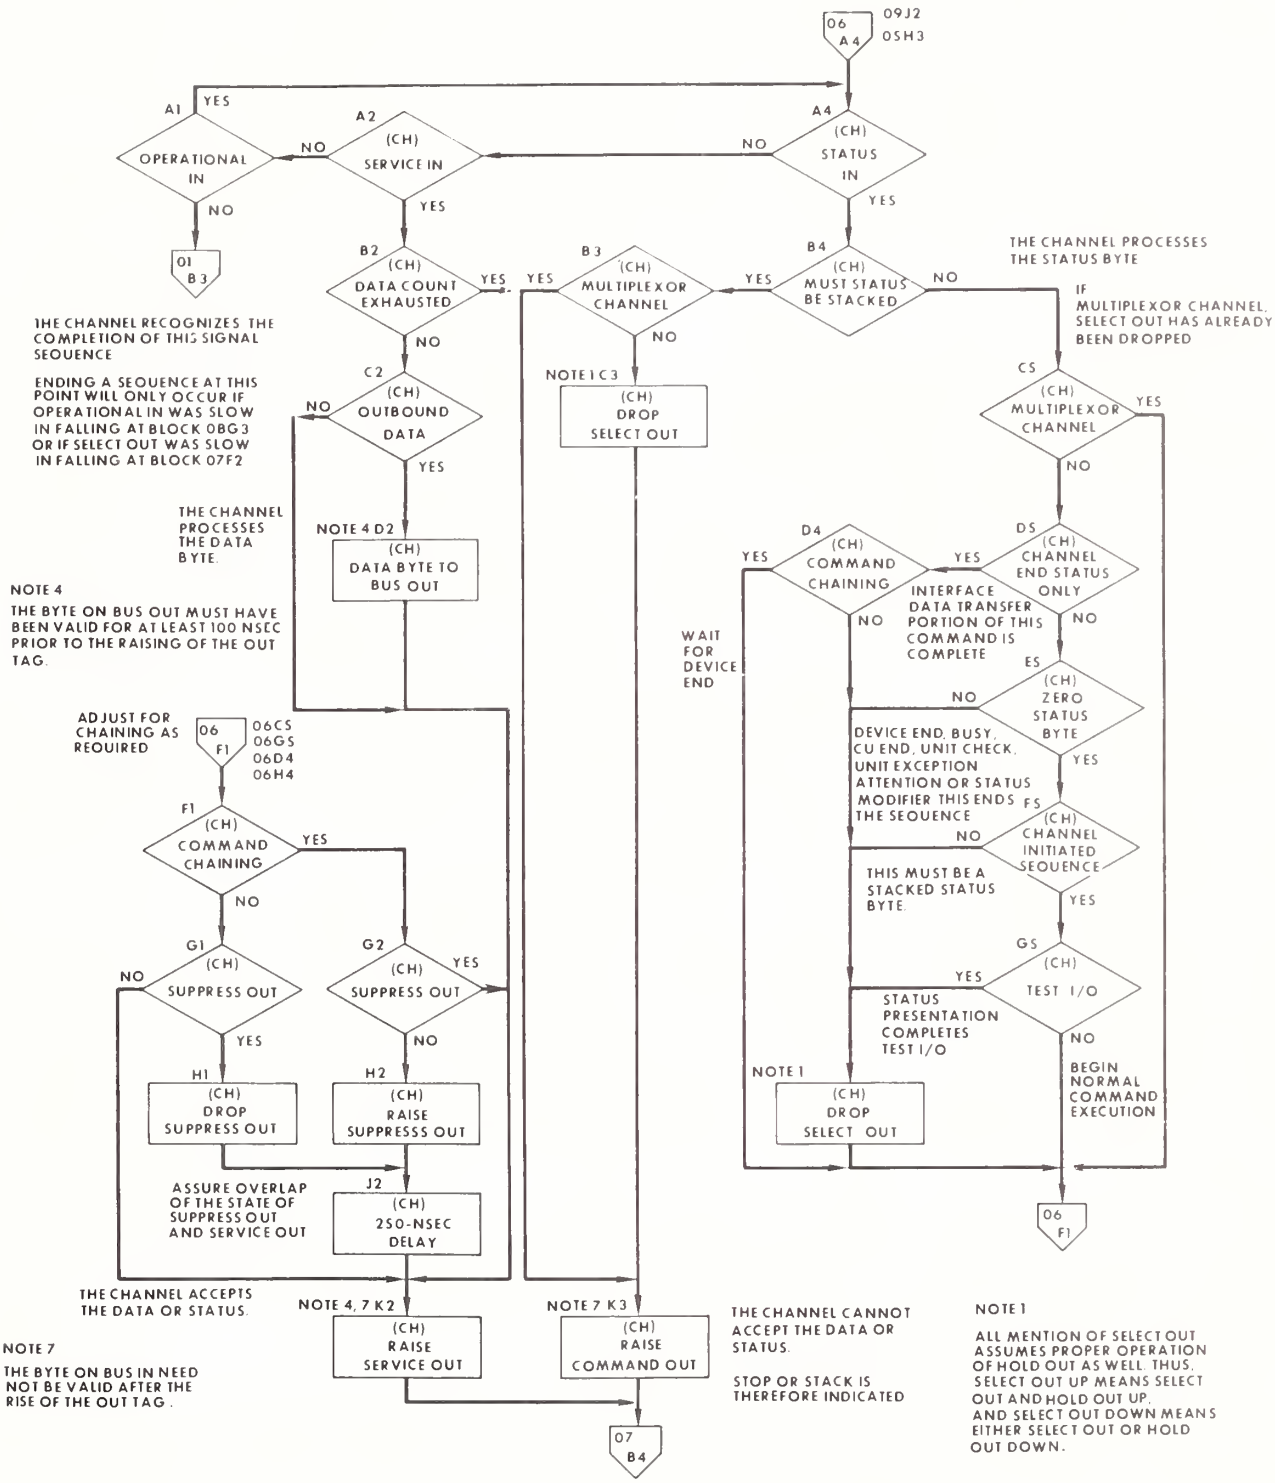

Flow Diagram 5. *Response to Status/Data Presentation*

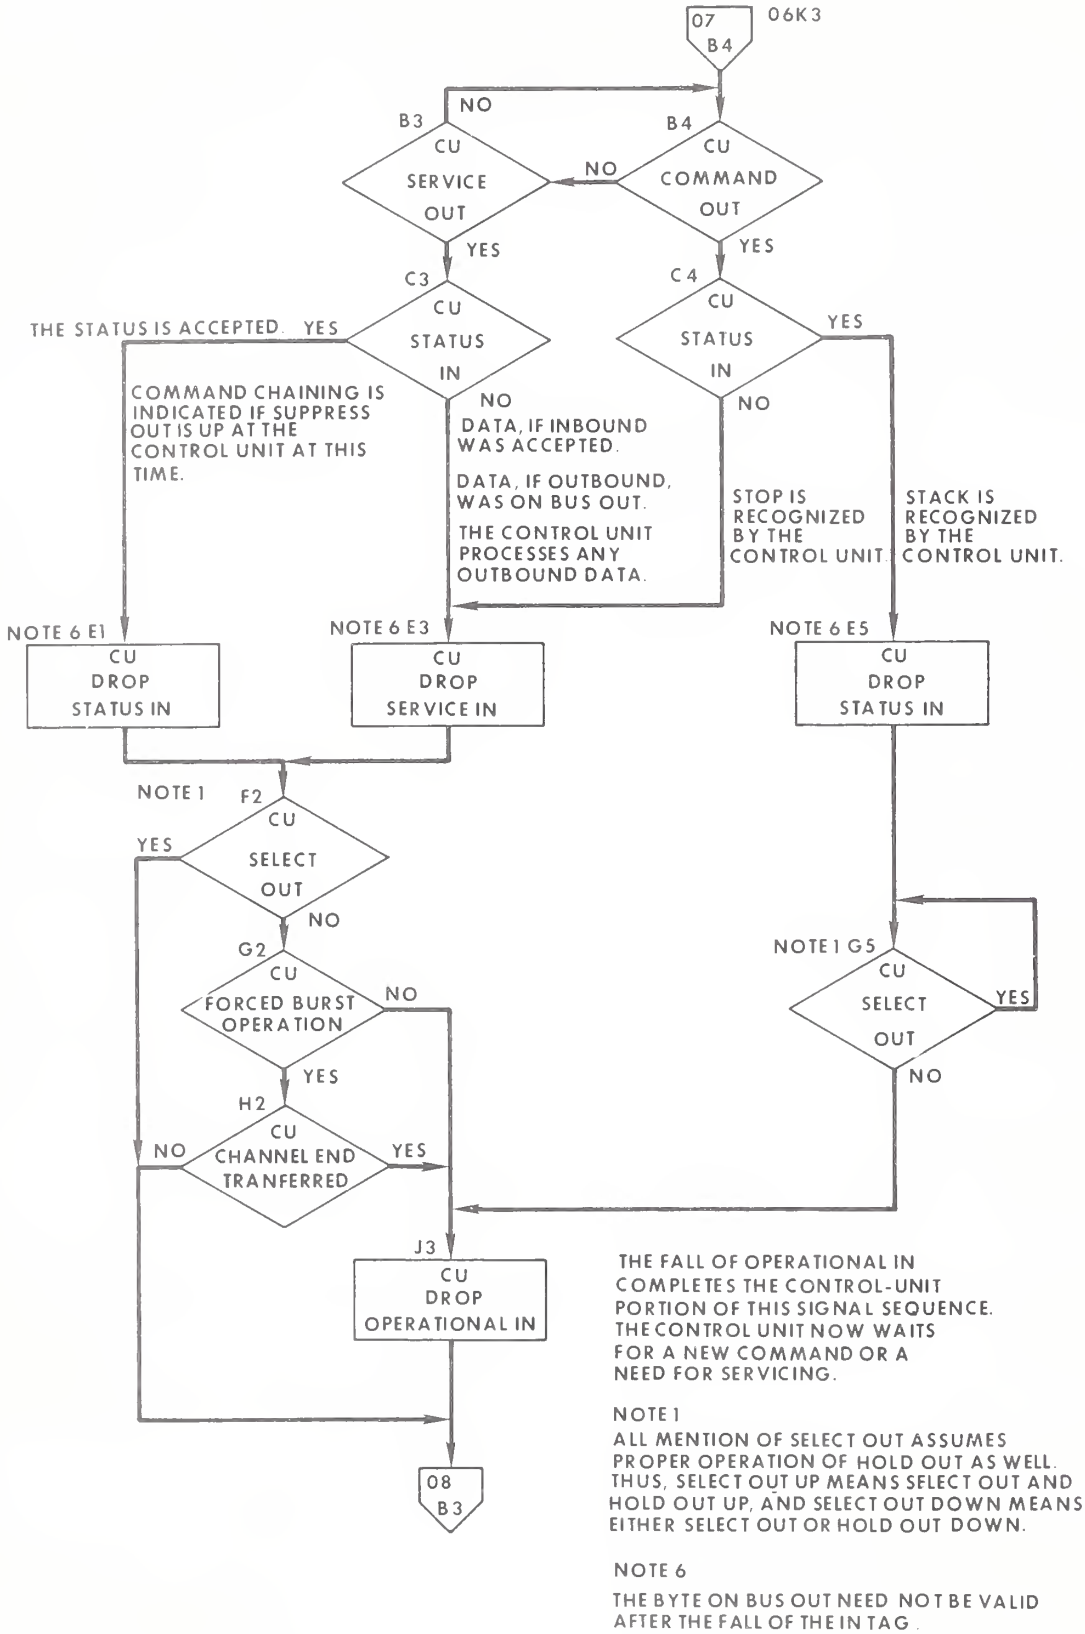

Flow Diagram 6. *Response to Stack/Stop Accept*

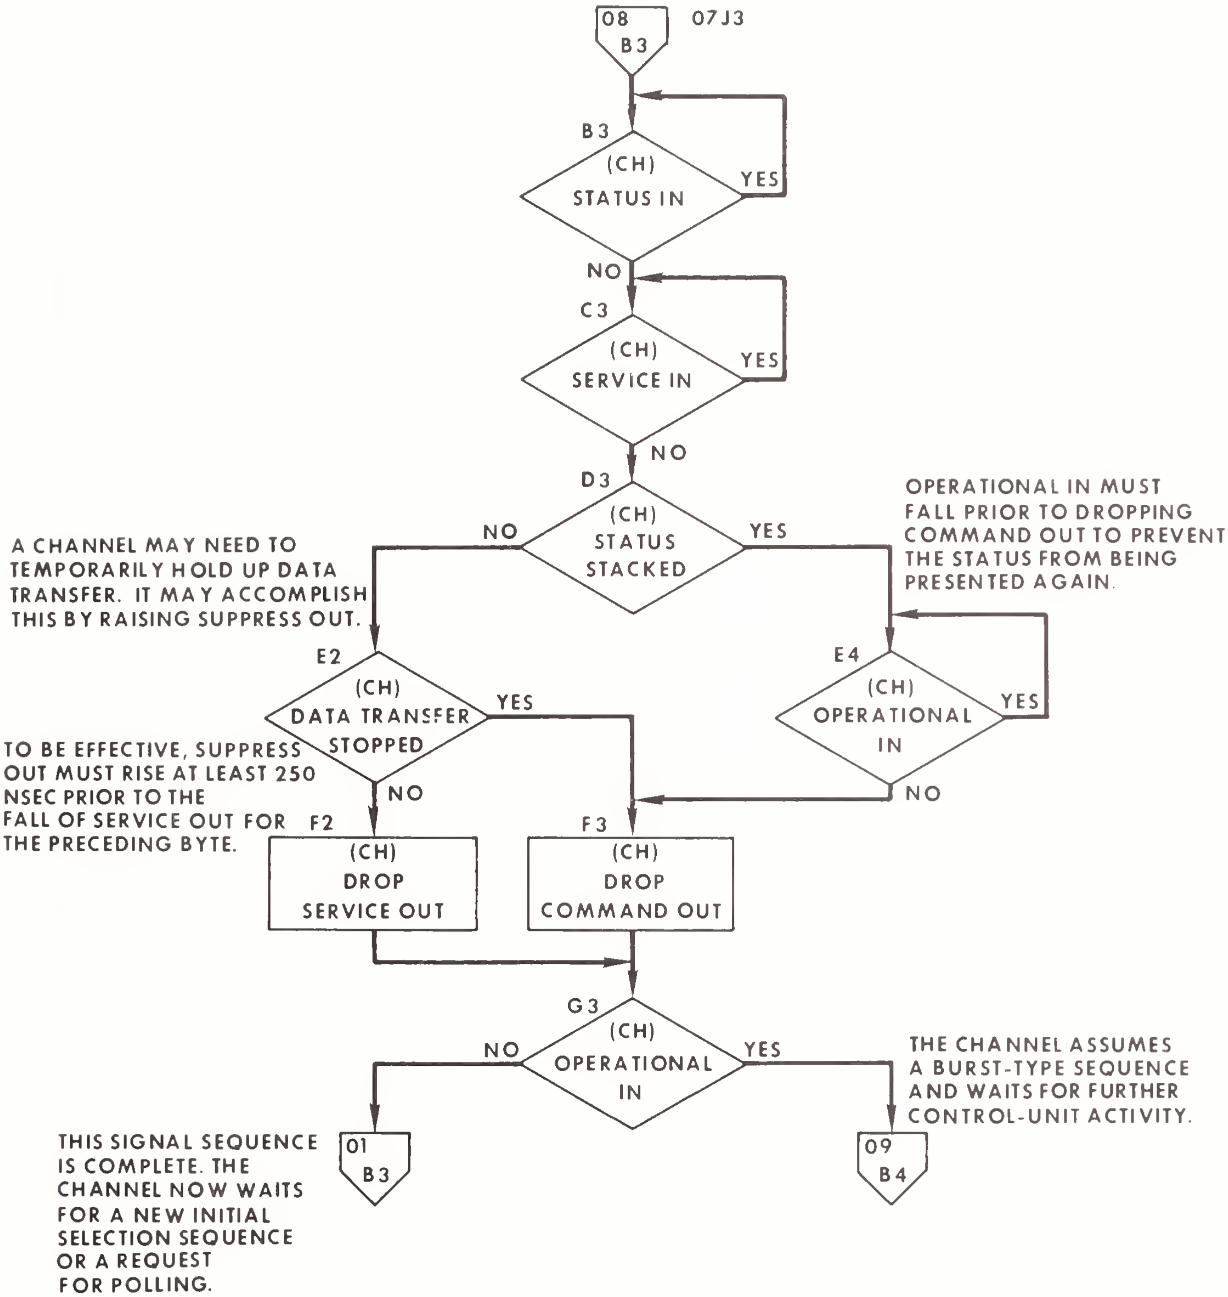

FLOW DIAGRAM 7. *Response to Fall of Status/Service in*

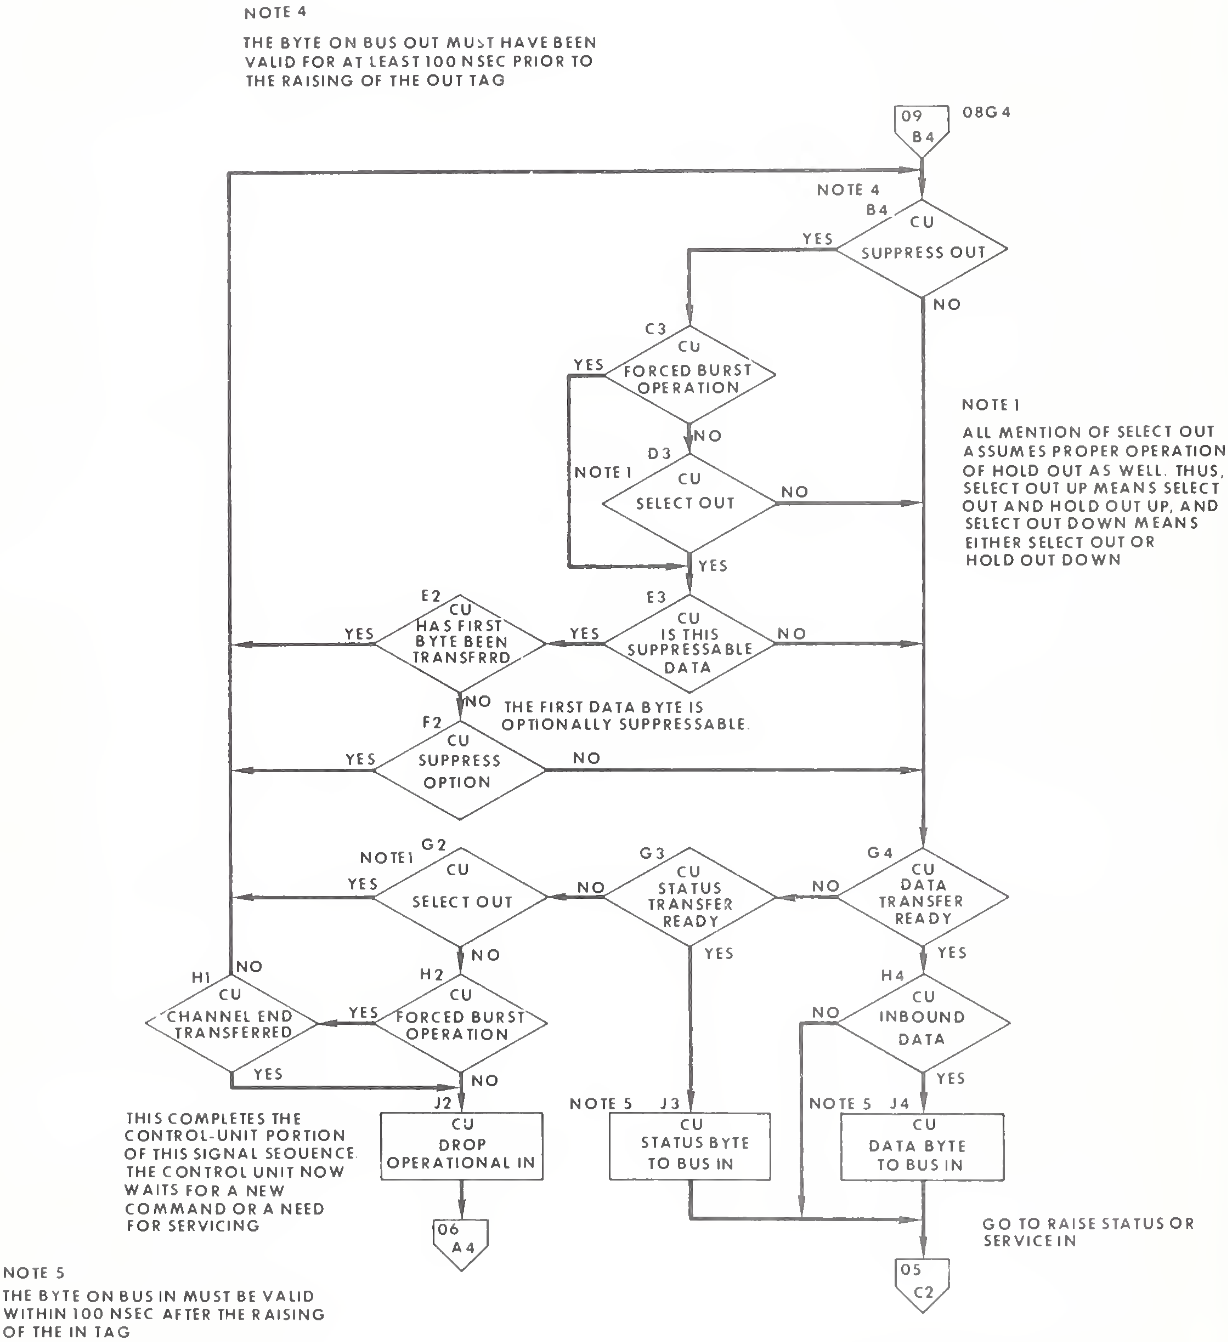

Flow Diagram 8. *Burst Mode Waiting Loop*

APPENDIX B.

SEQUENCE CHARTS

(This Appendix is not a part of the Standard, Specifications for I/O Channel Interface, but is included for information purposes only.)

Note: The terms "Byte Multiplexer Channel", "Block Multiplexer Channel", and "Selector Channel" are used to describe possible implementations of this standard.

U.S. GOVERNMENT PRINTING OFFICE: 1984 - 427-956 - 814/305

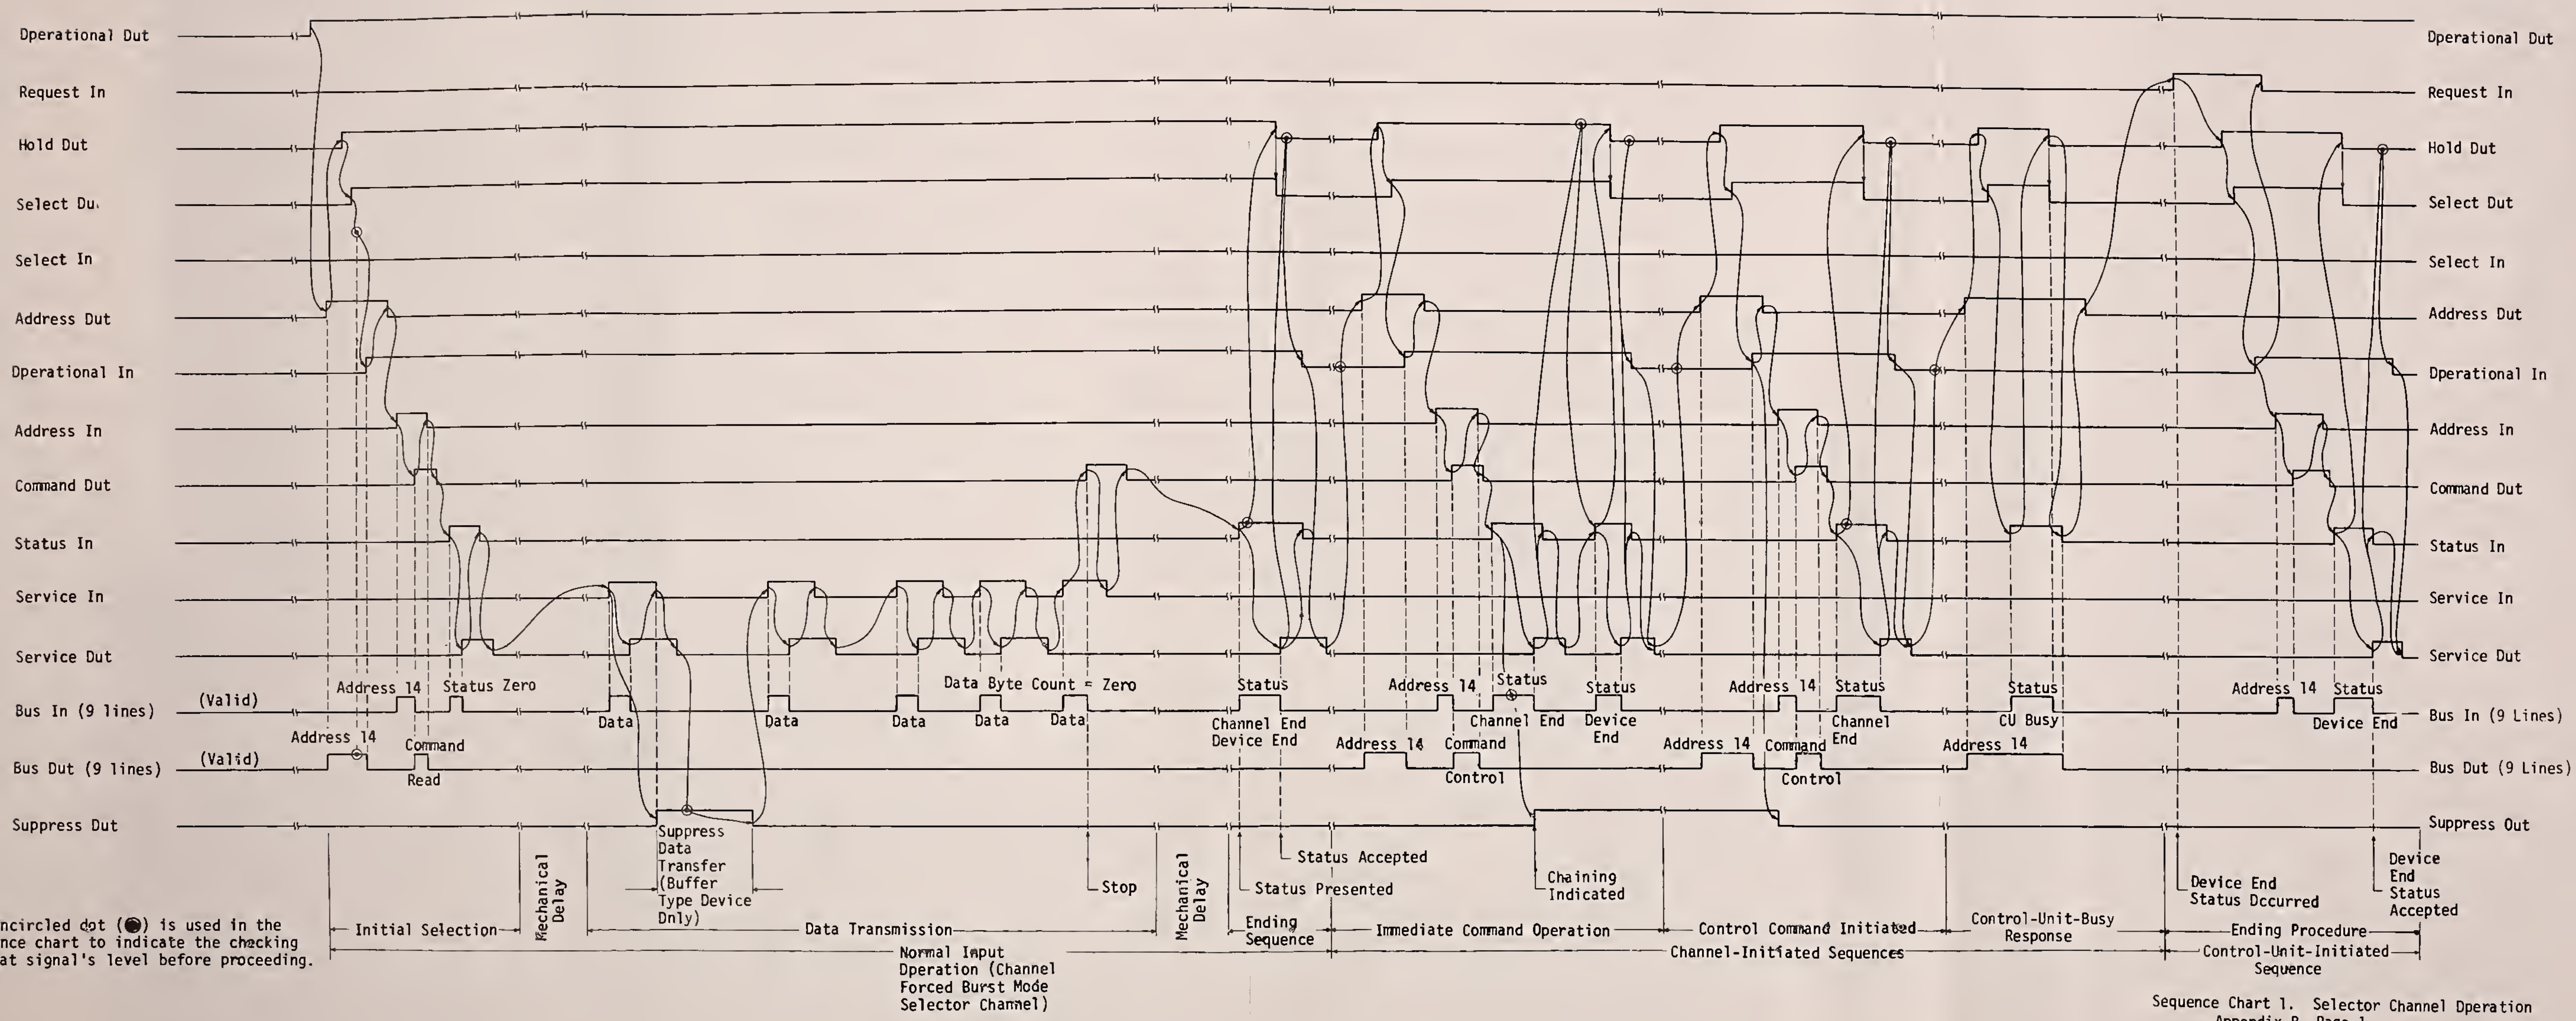

Note: The encircled dot (●) is used in the sequence chart to indicate the checking of that signal's level before proceeding.

Sequence Chart 1. Selector Channel Operation
Appendix B, Page 1

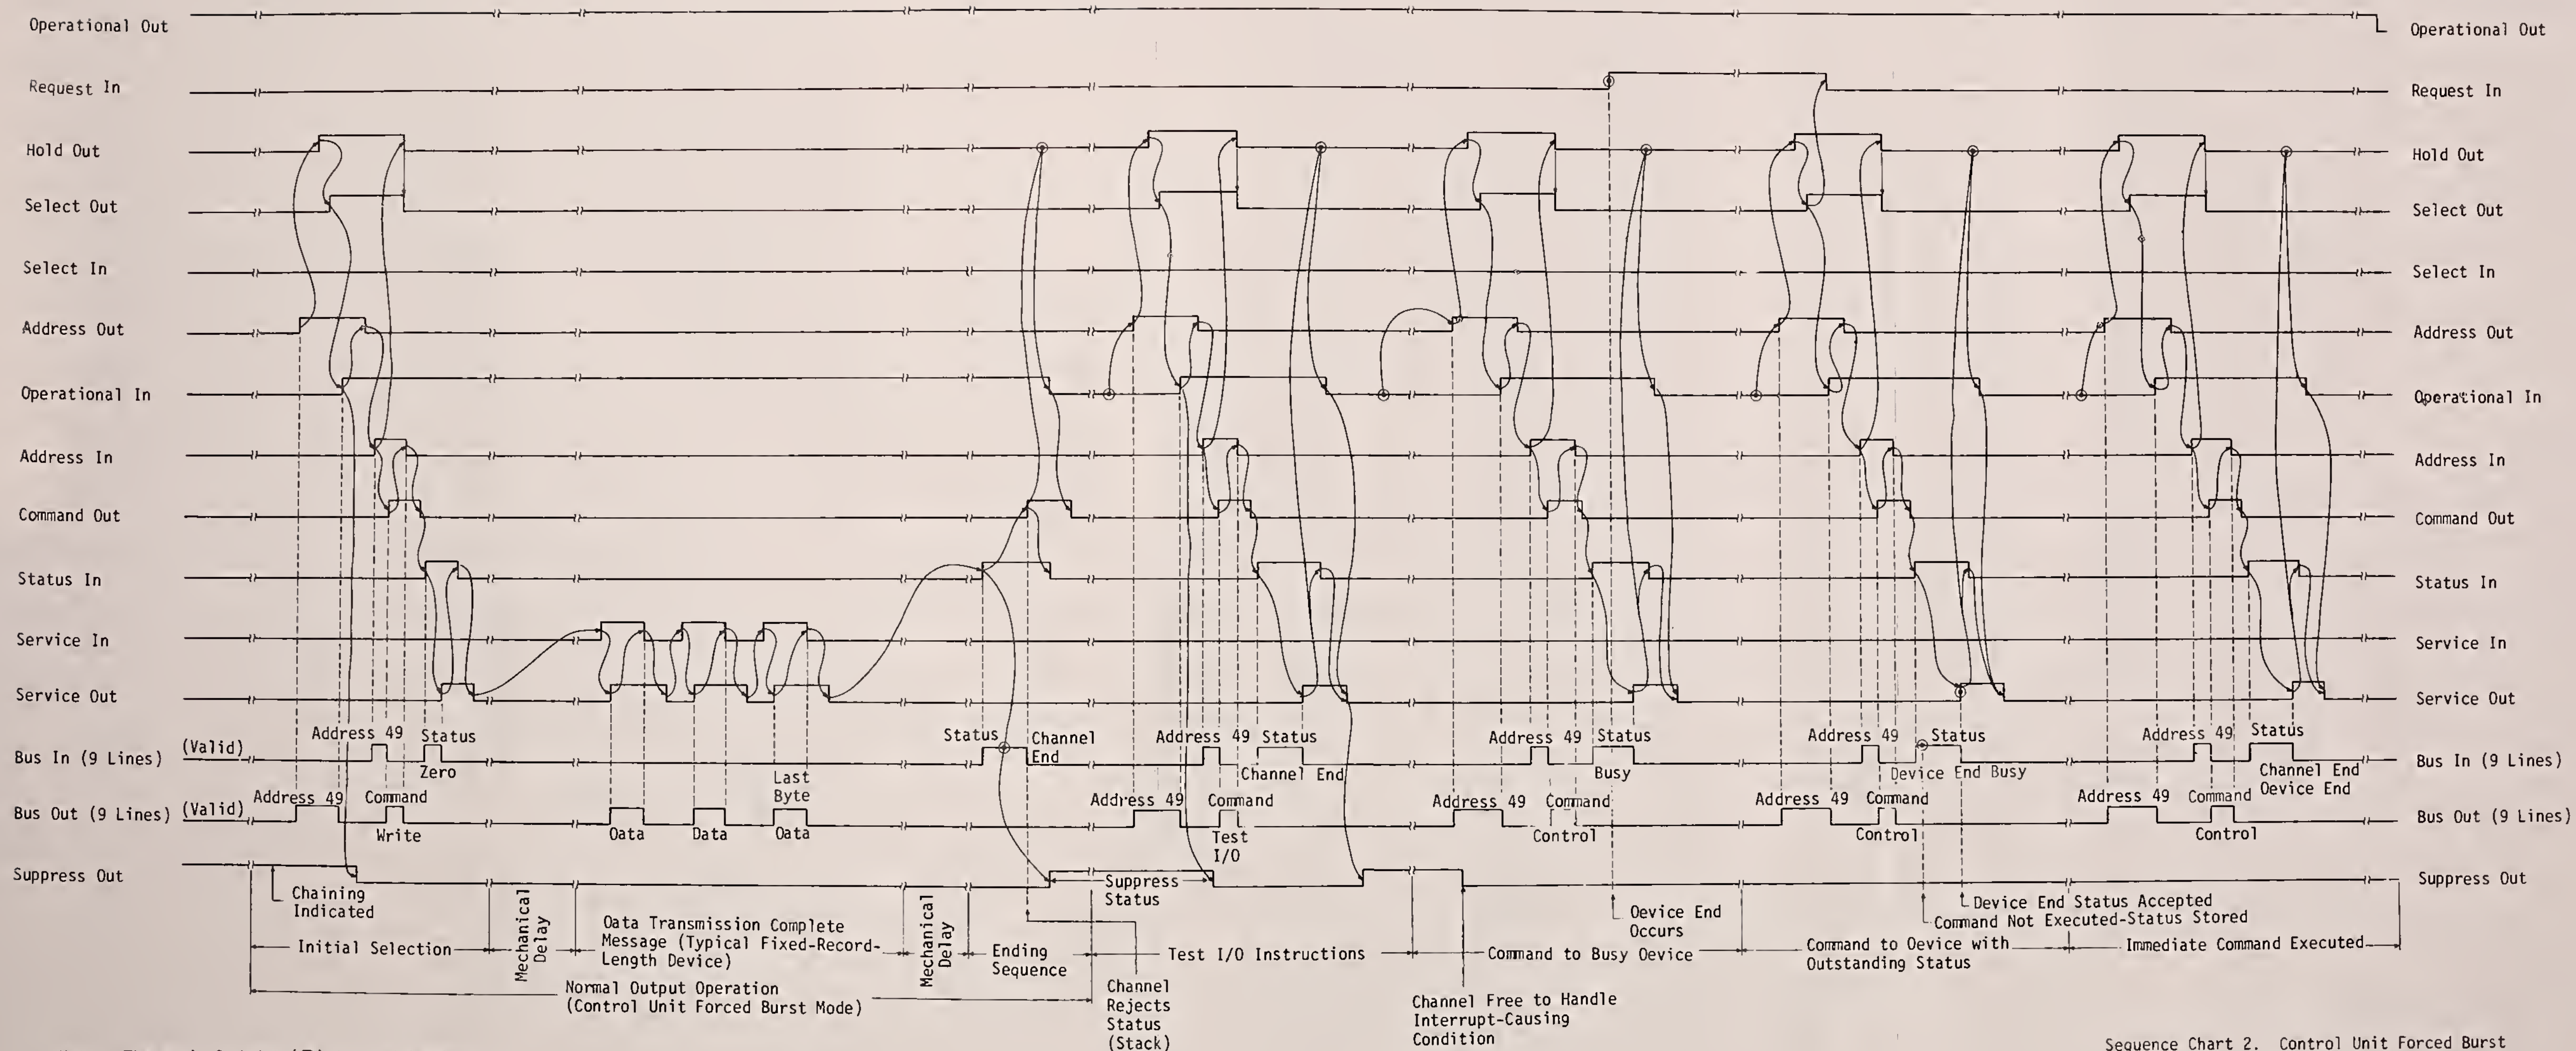

Note: The encircled dot (◉) is used in the sequence chart to indicate the checking of that signal's level before proceeding.

Sequence Chart 2. Control Unit Forced Burst Mode-Multiplexer Channel

Note: The encircled dot (◉) is used in the sequence chart to indicate the checking of that signal's level before proceeding.

Sequence Chart 3. Multiplexing Operation on Multiplexer Channel

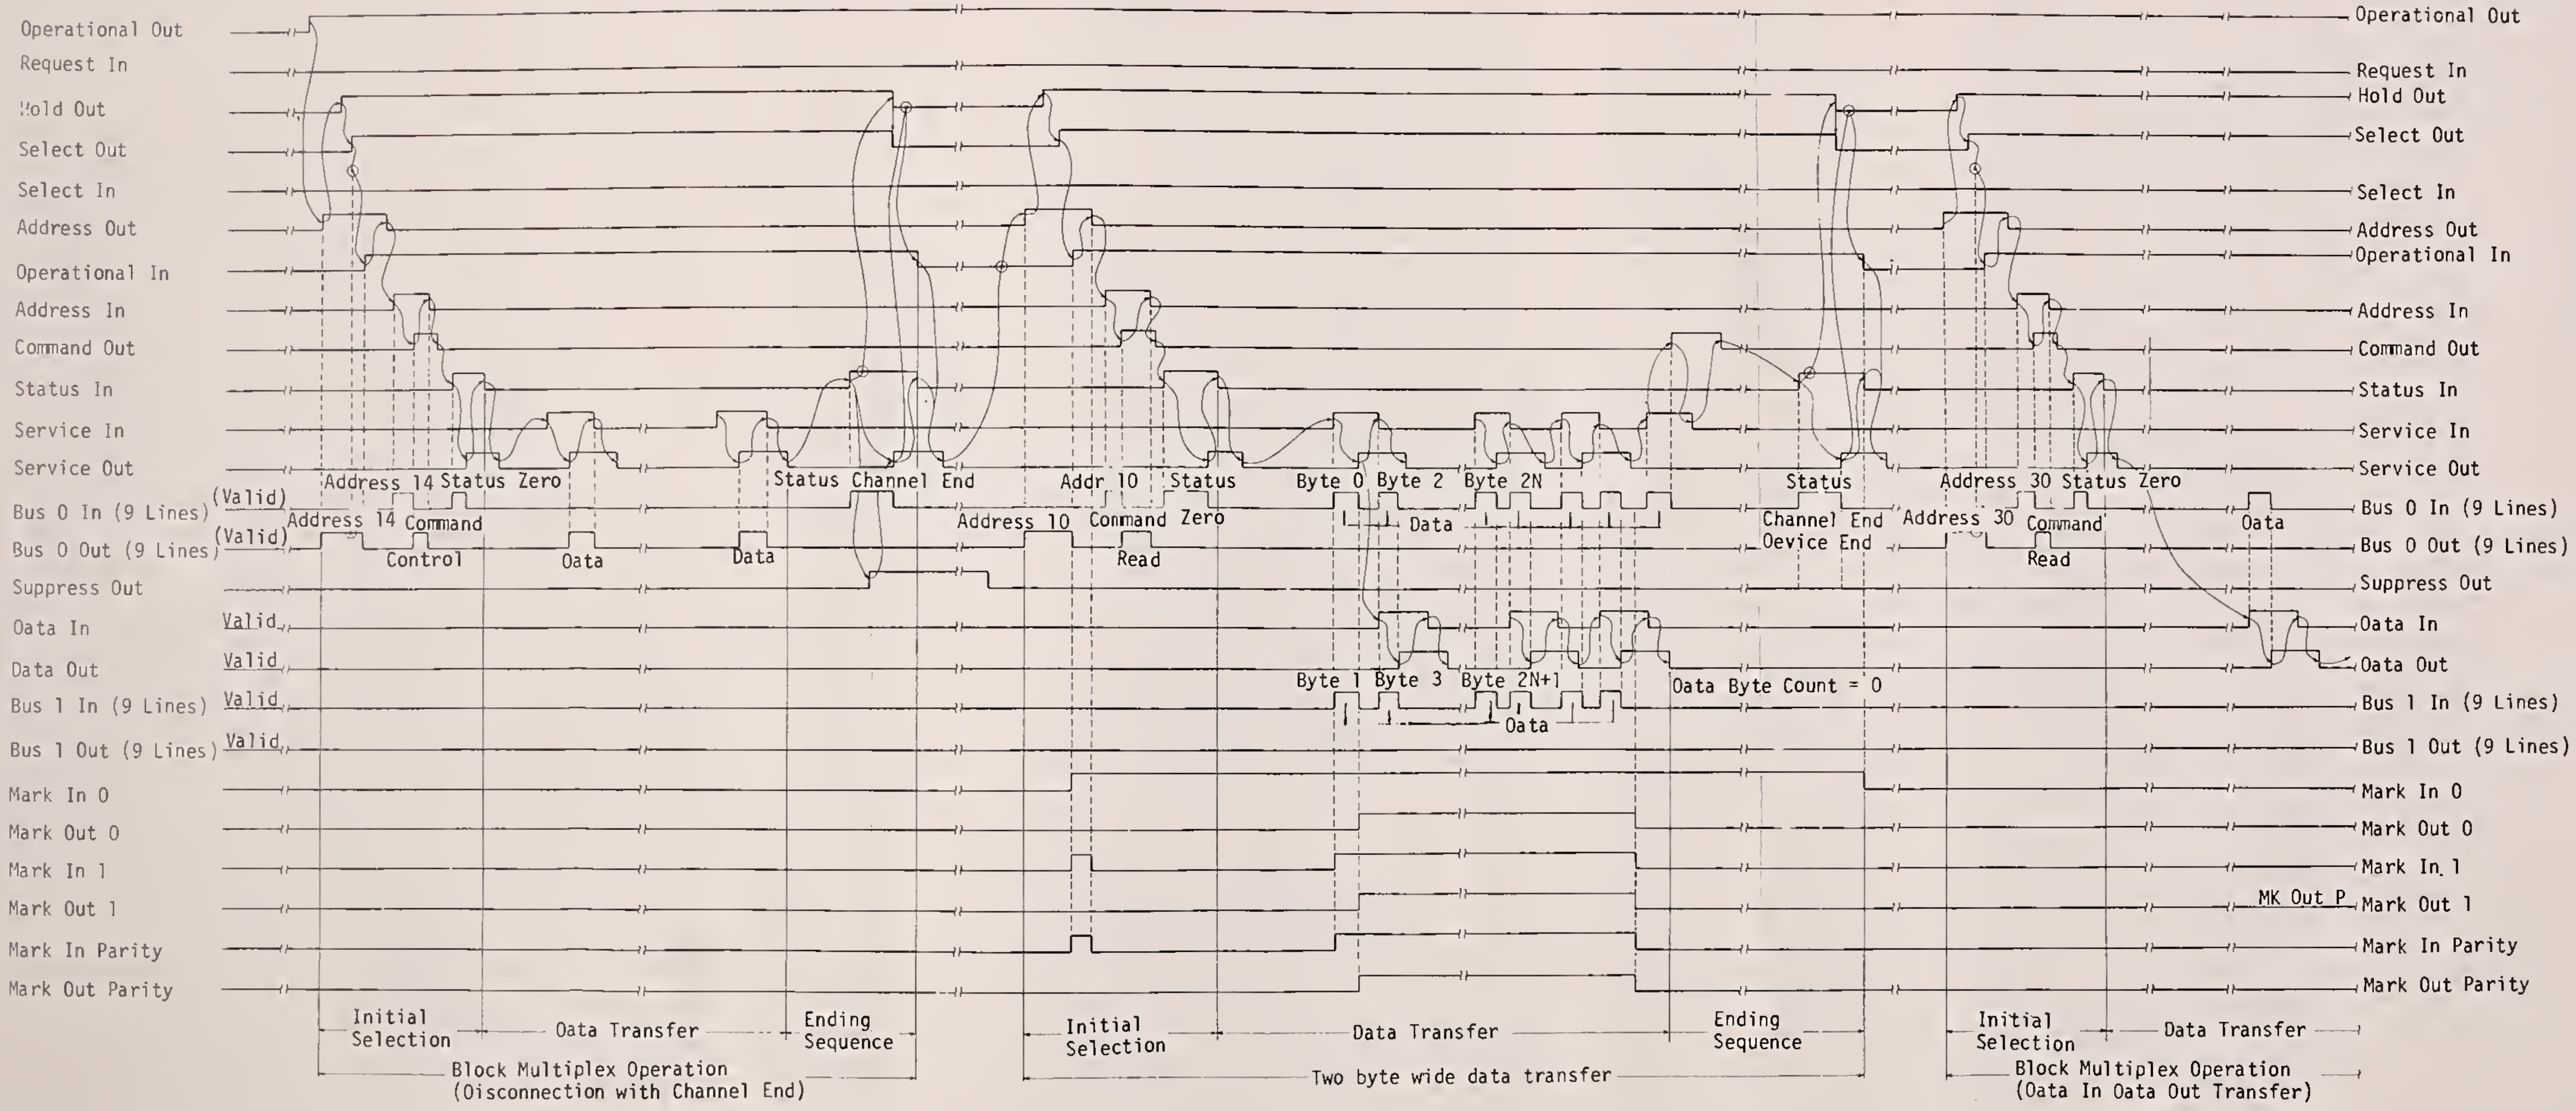

Sequence Chart 4. Block Multiplexer Channel Operation (Part 1 of 2)

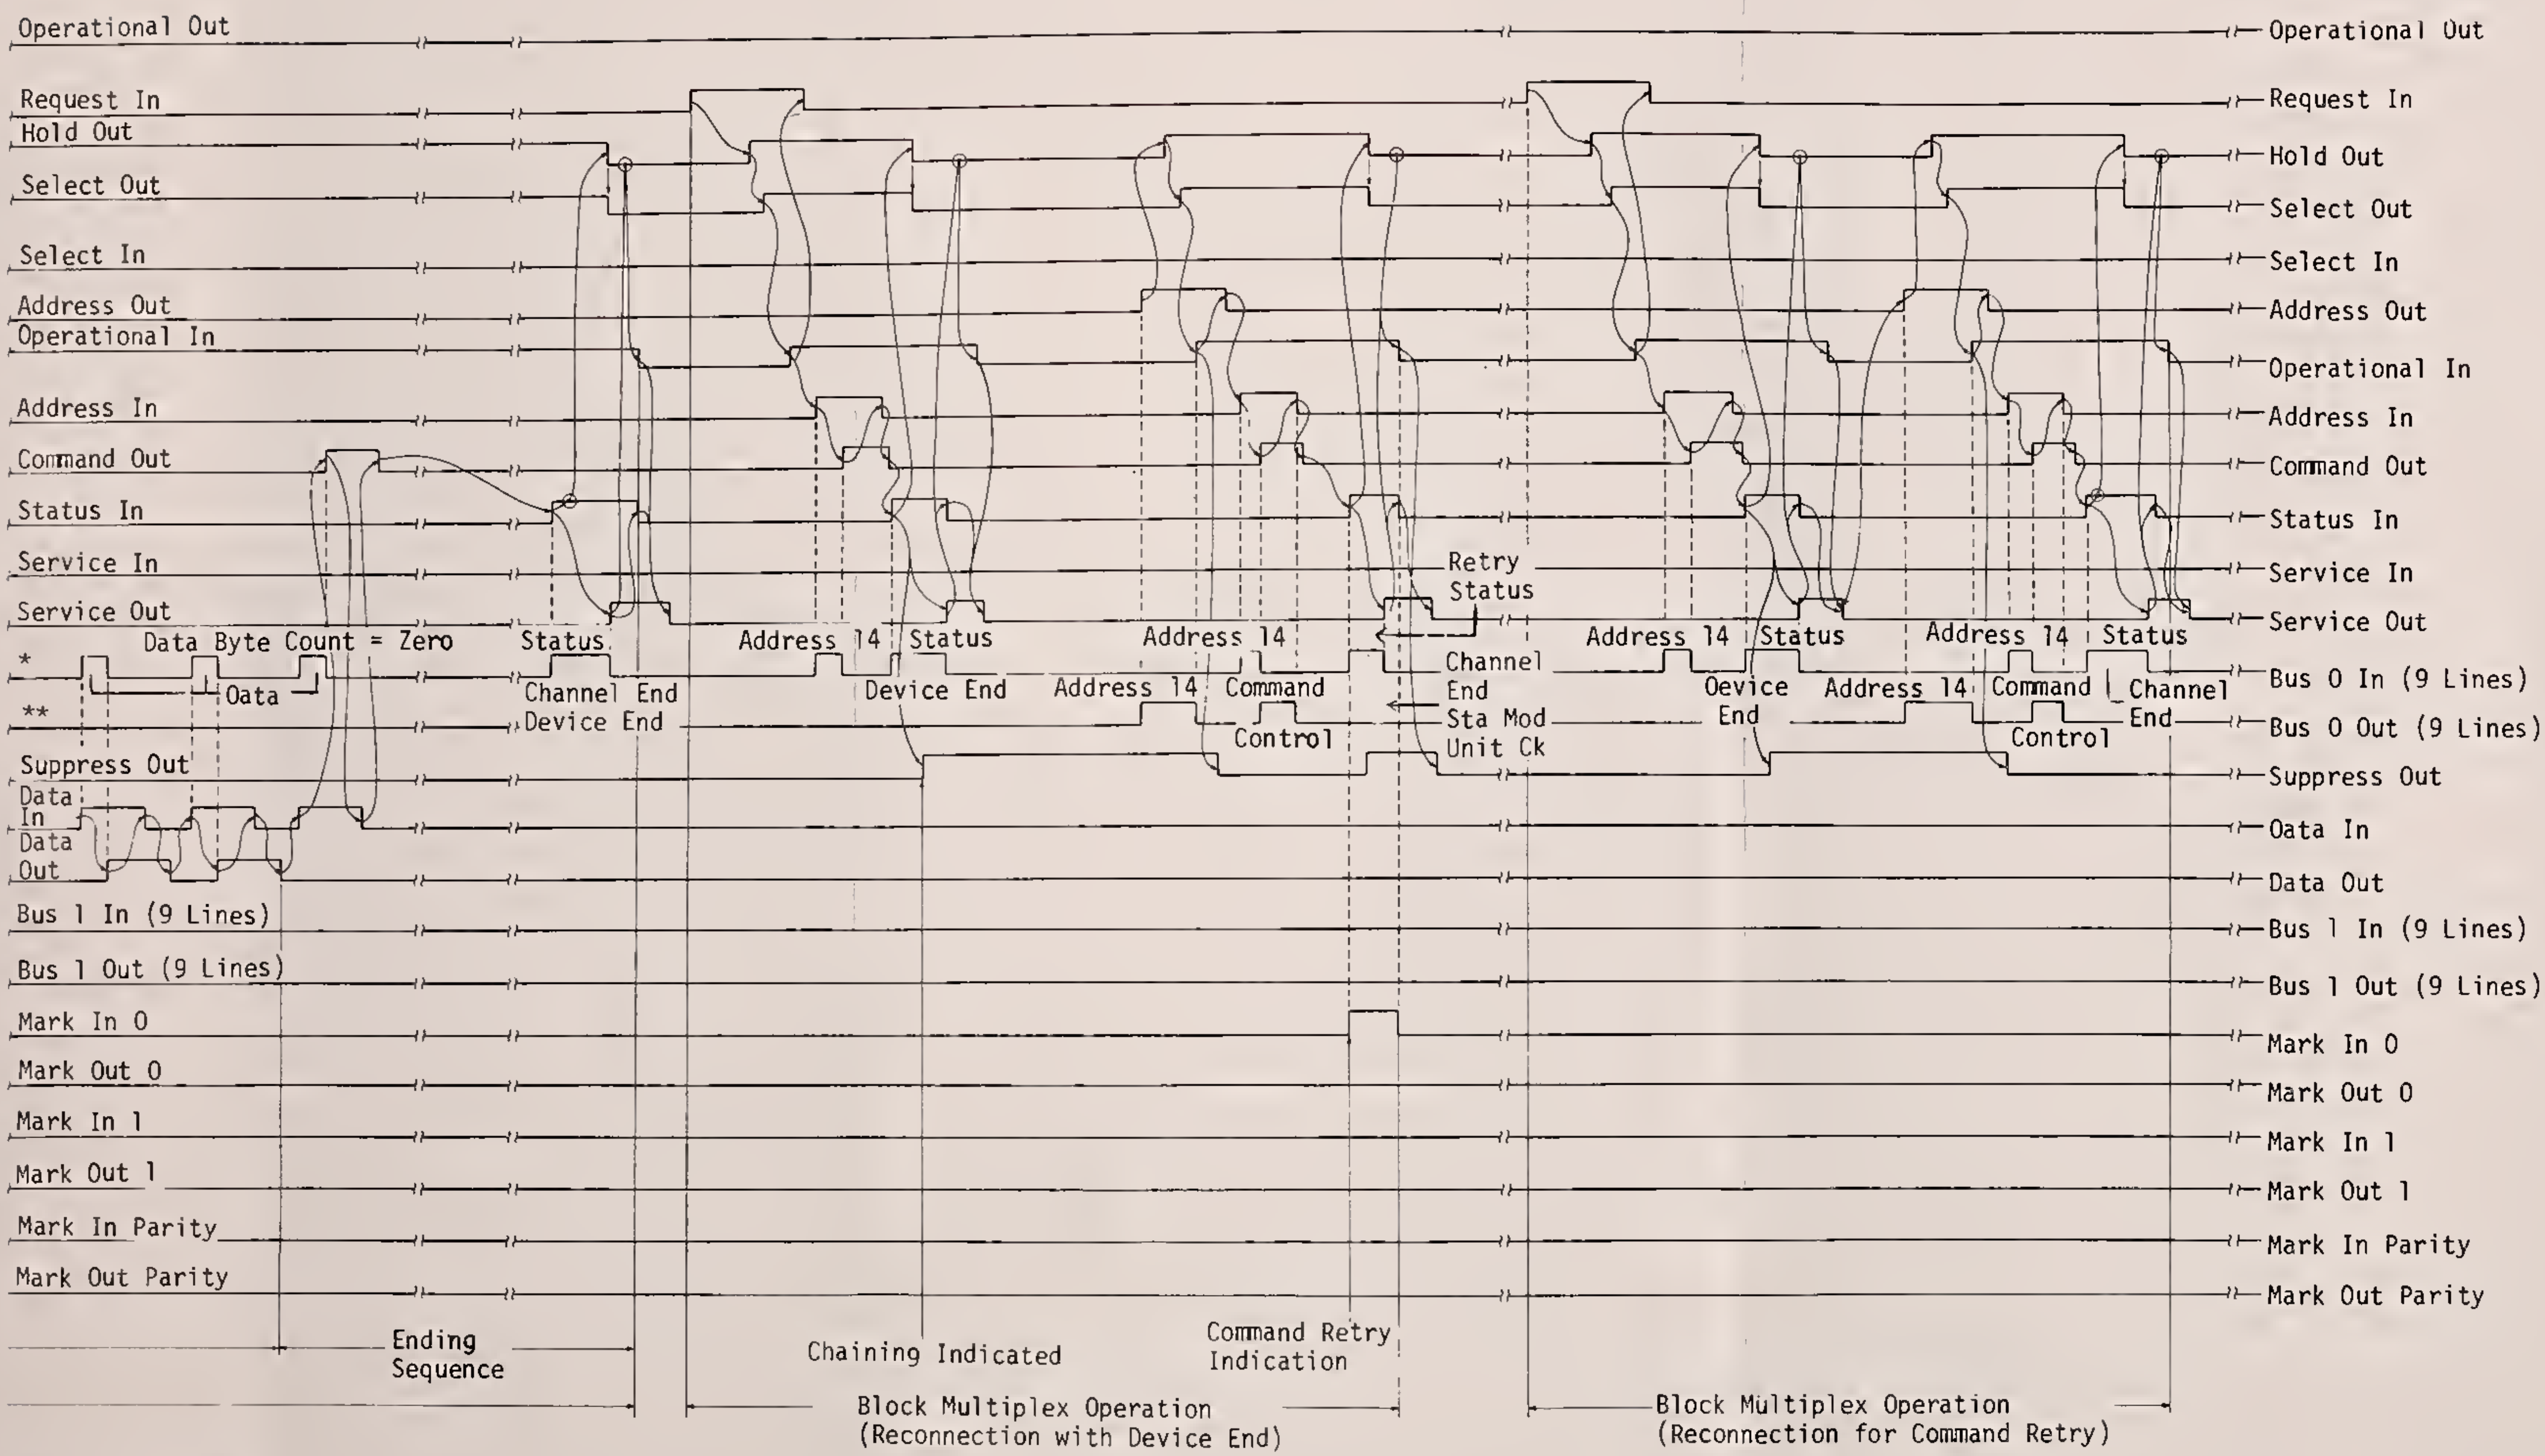

* Bus 0 In (9 Lines)

** Bus 0 Out (9 Lines)

Sequence Chart 4. Block Multiplexer Channel Operation (Part 2 of 2)

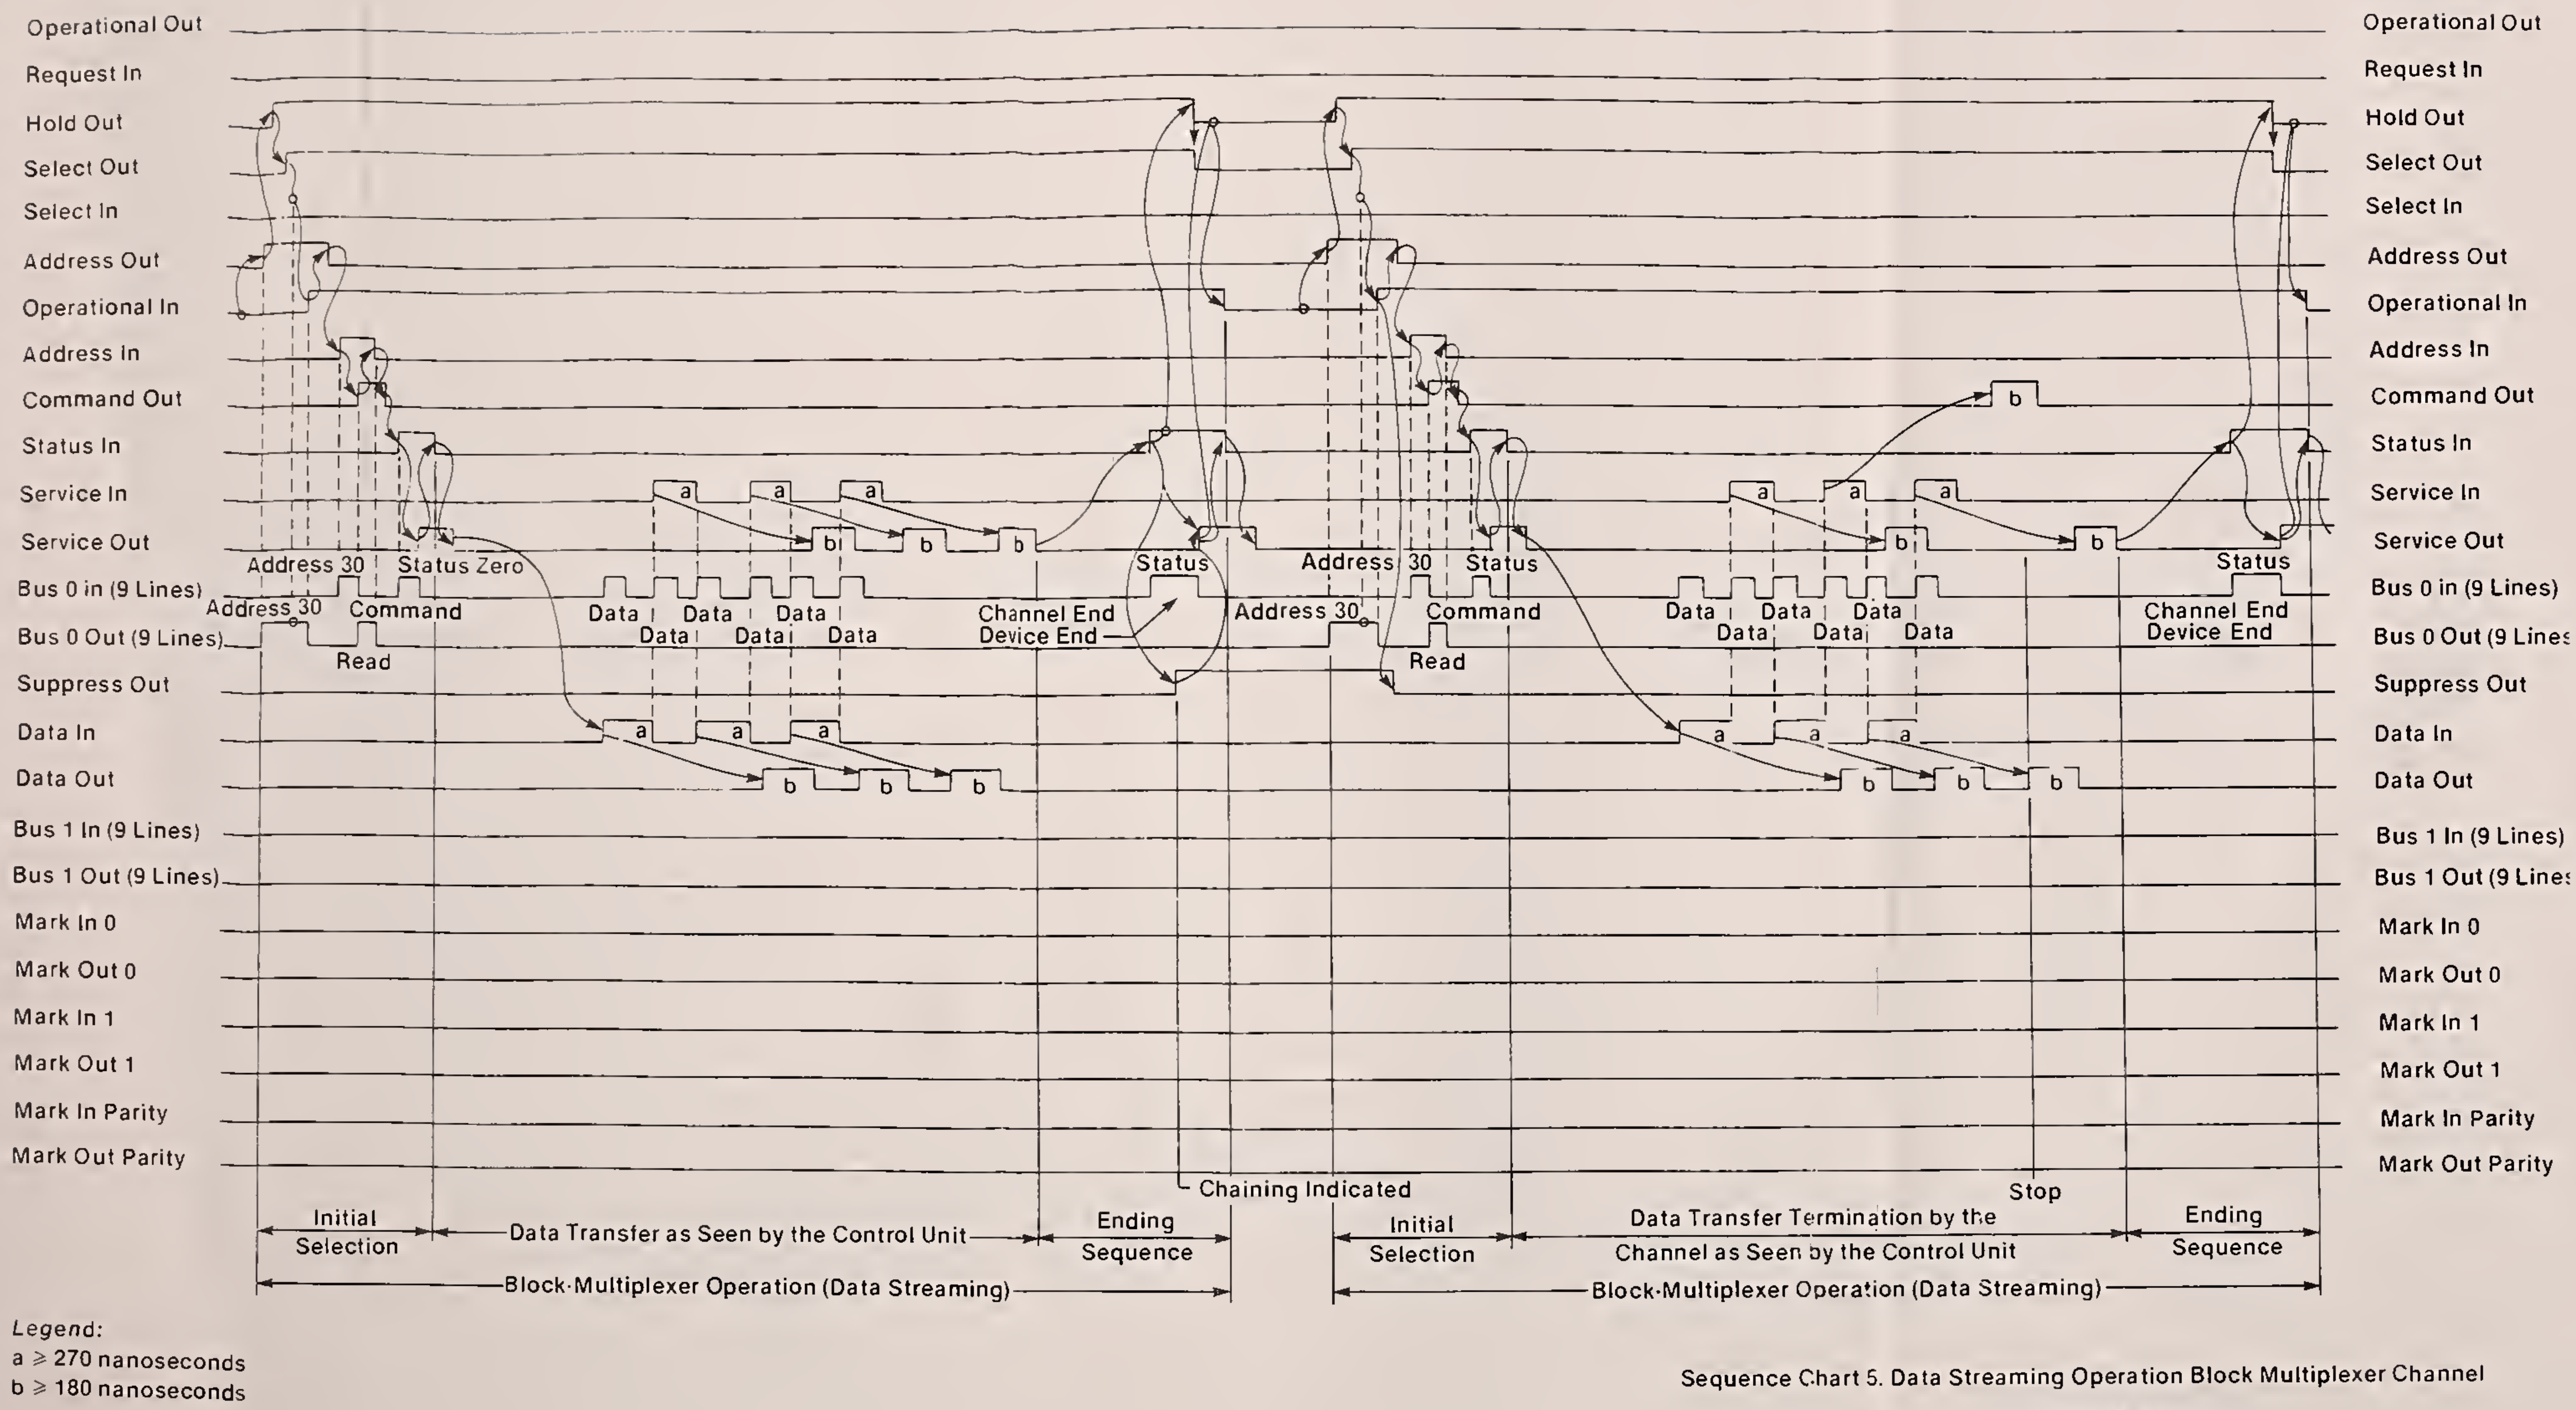

Sequence Chart 5. Data Streaming Operation Block Multiplexer Channel

NBS TECHNICAL PUBLICATIONS

PERIODICALS

JOURNAL OF RESEARCH—The Journal of Research of the National Bureau of Standards reports NBS research and development in those disciplines of the physical and engineering sciences in which the Bureau is active. These include physics, chemistry, engineering, mathematics, and computer sciences. Papers cover a broad range of subjects, with major emphasis on measurement methodology and the basic technology underlying standardization. Also included from time to time are survey articles on topics closely related to the Bureau's technical and scientific programs. As a special service to subscribers each issue contains complete citations to all recent Bureau publications in both NBS and non-NBS media. Issued six times a year. Annual subscription: domestic $18; foreign $22.50. Single copy, $5.50 domestic; $6.90 foreign.

NONPERIODICALS

Monographs—Major contributions to the technical literature on various subjects related to the Bureau's scientific and technical activities.

Handbooks—Recommended codes of engineering and industrial practice (including safety codes) developed in cooperation with interested industries, professional organizations, and regulatory bodies.

Special Publications—Include proceedings of conferences sponsored by NBS, NBS annual reports, and other special publications appropriate to this grouping such as wall charts, pocket cards, and bibliographies.

Applied Mathematics Series—Mathematical tables, manuals, and studies of special interest to physicists, engineers, chemists, biologists, mathematicians, computer programmers, and others engaged in scientific and technical work.

National Standard Reference Data Series—Provides quantitative data on the physical and chemical properties of materials, compiled from the world's literature and critically evaluated. Developed under a worldwide program coordinated by NBS under the authority of the National Standard Data Act (Public Law 90-396).

NOTE: The principal publication outlet for the foregoing data is the Journal of Physical and Chemical Reference Data (JPCRD) published quarterly for NBS by the American Chemical Society (ACS) and the American Institute of Physics (AIP). Subscriptions, reprints, and supplements available from ACS, 1155 Sixteenth St., NW, Washington, DC 20056.

Building Science Series—Disseminates technical information developed at the Bureau on building materials, components, systems, and whole structures. The series presents research results, test methods, and performance criteria related to the structural and environmental functions and the durability and safety characteristics of building elements and systems.

Technical Notes—Studies or reports which are complete in themselves but restrictive in their treatment of a subject. Analogous to monographs but not so comprehensive in scope or definitive in treatment of the subject area. Often serve as a vehicle for final reports of work performed at NBS under the sponsorship of other government agencies.

Voluntary Product Standards—Developed under procedures published by the Department of Commerce in Part 10, Title 15, of the Code of Federal Regulations. The standards establish nationally recognized requirements for products, and provide all concerned interests with a basis for common understanding of the characteristics of the products. NBS administers this program as a supplement to the activities of the private sector standardizing organizations.

Consumer Information Series—Practical information, based on NBS research and experience, covering areas of interest to the consumer. Easily understandable language and illustrations provide useful background knowledge for shopping in today's technological marketplace.

Order the **above** *NBS publications from: Superintendent of Documents, Government Printing Office, Washington, DC 20402.*

Order the **following** *NBS publications—FIPS and NBSIR's—from the National Technical Information Service, Springfield, VA 22161.*

Federal Information Processing Standards Publications (FIPS PUB)—Publications in this series collectively constitute the Federal Information Processing Standards Register. The Register serves as the official source of information in the Federal Government regarding standards issued by NBS pursuant to the Federal Property and Administrative Services Act of 1949 as amended, Public Law 89-306 (79 Stat. 1127), and as implemented by Executive Order 11717 (38 FR 12315, dated May 11, 1973) and Part 6 of Title 15 CFR (Code of Federal Regulations).

NBS Interagency Reports (NBSIR)—A special series of interim or final reports on work performed by NBS for outside sponsors (both government and non-government). In general, initial distribution is handled by the sponsor; public distribution is by the National Technical Information Service, Springfield, VA 22161, in paper copy or microfiche form.

U.S. DEPARTMENT OF COMMERCE
National Technical Information Service
5285 Port Royal Road
Springfield, Virginia 22161

OFFICIAL BUSINESS

POSTAGE AND FEES PAID
U.S. DEPARTMENT OF COMMERCE
COM-211

3rd Class Bulk Rate

U.S. MAIL

12-18-90
Vol. 55 No. 243
Pages 51895-52036

federal register

Tuesday
December 18, 1990

National Institute of Standards and Technology

[Docket No. 900101-0219]

RIN 0693-AA59

Approval of Revisions to Federal Information Processing Standards (FIPS) Family of Input/Output Interface Standards

AGENCY: National Institute of Standards and Technology (NIST), Commerce.

ACTION: The purpose of this notice is to announce that the Secretary of Commerce has approved revisions to the Federal Information Processing Standards (FIPS) family of input/output interface standards, and has approved discontinuation of the exclusion and verification lists for these standards.

SUMMARY: On March 20, 1990, notice was published in the Federal Register (55 FR 10272) proposing revision of Federal Information Processing Standards (FIPS) 60-2, 61-1, 62, 63-1, 97, 111, 130, and 131 to make them non-mandatory, and discontinue the exclusion and verification lists for these standards. This proposal superseded the proposal for revision of these standards announced in the Federal Register (52 FR 44462) of November 19, 1987. Procedures for the Exclusion List for FIPS 60, 61, 62, 63, and 97 were published in the Federal Register on

September 3, 1982 (47 FR 38959–38960). Procedures for the Verification List for FIPS 60, 61, 62, 63, and 97 were published in the Federal Register on December 11, 1979 (44 FR 71444–71445) and on April 7, 1981 (46 FR 20719–20720).

The written comments submitted by interested parties and other material available to the Department relevant to these proposed revisions were reviewed by NIST. On the basis of this review, NIST recommended that the Secretary approve revisions to the input/output family of standards and approve discontinuation of the exclusion and verification lists for these standards. NIST prepared a detailed justification document for the Secretary's review in support of those recommendations.

This notice provides only the changes to the revised standards.

EFFECTIVE DATE: These revisions are effective December 18, 1990.

ADDRESSES: Interested parties may obtain copies of FIPS PUBS 60–2, 61–1, 62, 63–1, 97, 111, 130, and 131 from the National Technical Information Service, U.S. Department of Commerce, Springfield, VA 22161.

FOR FURTHER INFORMATION CONTACT: Ms. Shirley Radack, National Institute of Standards and Technology, Gaithersburg, MD 20899, telephone (301) 975–2833.

SUPPLEMENTARY INFORMATION: Under the provisions of 40 U.S.C. 759(d), the Secretary of Commerce is authorized to promulgate standards and guidelines for Federal computer systems, and to make such standards compulsory and binding to the extent to which the Secretary determines necessary to improve the efficiency of operation, or security and privacy of Federal computer systems.

The family of I/O interface standards currently includes:

a. FIPS 60–2, I/O Channel Interface, revised July 29, 1983.

b. FIPS 61–1, Channel Level Power Control Interface, revised July 13, 1982.

c. FIPS 62, Operational Specifications for Magnetic Tape Subsystems, revised December 30, 1980.

d. FIPS 63–1, Operational Specifications for Variable Block Rotating Mass Storage Subsystems, revised April 14, 1983; Supplement to FIPS PUB. 63–1, Additional Operational Specifications for Variable Block Rotating Mass Storage Subsystems, April 14, 1983.

e. FIPS 97, Operational Specifications for Fixed Block Rotating Mass Storage Subsystems, February 4, 1983.

f. FIPS 111, Storage Module Interfaces (with extensions for enhanced storage module interfaces), April 18, 1985.

g. FIPS 130, Intelligent Peripheral Interface (IPI), July 16, 1987.

h. FIPS 131, Small Computer System Interface (SCSI) July 16, 1987.

The following revisions are being made effective immediately upon publication. A delayed effective date is not required because these standards are exempt from the Administrative Procedure Act by U.S.C. 553(a)(2).

Revisions to Federal Information Processing Standards 60–2, 61–1, 62, 63–1, 97, 111, 130, and 131.

FIPS 60–2, I/O Channel Interface, is revised as follows:

Applicability. This standard addresses the interconnection of computer peripheral equipment as a part of ADP systems for the following types of peripherals: (1) Magnetic tape equipment employing open reel-to-reel magnetic tape storage devices, specifically excluding magnetic tape cassette and tape cartridge storage devices, (2) magnetic disk storage equipment employing disk drives each having a capacity greater than 7 megabytes per storage module, excluding flexible disk and disk cartridge devices having a smaller storage capacity per device, and (3) other peripheral equipment employing peripheral device types for which operational specifications standards have been issued as Federal Information Processing Standards. This standard is recommended for use in the acquisition of peripheral equipment for ADP systems with input/output channel interfaces as specified in the technical specifications, when it is determined that interchange of equipment between different systems is likely.

Implementation. The original version of this standard became effective December 13, 1979. The first revision became effective June 23, 1980, and the second revision became effective July 29, 1983. This revision becomes effective December 18, 1990.

Waivers. This standard is non-mandatory. No waivers are required.

FIPS 61–1, Channel Level Power Control Interface, is revised as follows:

Applicability. This standard addresses the power control interface in connecting computer peripheral equipment to ADP systems. It is recommended for use w/ when FIPS 60–2 is used, when it is determined that interchange of equipment between different systems is likely.

Implementation. The original version of this standard became effective June 23, 1980, and the first revision became effective July 13, 1982. This revision becomes effective December 18, 1990.

Waivers. This standard is non-mandatory. No waivers are required.

FIPS 62, Operational Specifications for Magnetic Tape Subsystems, is revised as follows:

Applicability. This standard addresses magnetic tape equipment connected to ADP systems through FIPS 60 interfaces. It is recommended for use in the acquisition of such equipment, when it is determined that interchange of equipment between different systems is likely.

Implementation. The original version of this standard became effective June 23, 1980. This revision becomes effective December 18, 1990.

Waivers. This standard is non-mandatory. No waivers are required.

FIPS 63–1, Operational Specifications for Variable Block Rotating Mass Storage Subsystems, is revised as follows:

Applicability. This standard addresses peripheral device dependent operational interfaces for connecting variable block rotating mass storage equipment to ADP systems through FIPS 60 interfaces. It is recommended for use in the acquisition of such variable block rotating mass storage equipment for connection to ADP systems, when it is determined that interchange of equipment between different systems is likely.

Implementation. This standard became effective June 23, 1980, and the first revision became effective April 14, 1983. This revision becomes effective December 18, 1990.

Waivers. This standard is non-mandatory. No waivers are required.

FIPS 97, Operational Specifications for Fixed Block Rotating Mass Storage Subsystems, is revised as follows:

Applicability. This standard addresses the peripheral device dependent operational interface specifications for connecting fixed block rotating mass storage equipment to ADP systems through FIPS 60 interfaces. It is recommended for use in the acquisition of such fixed block rotating mass storage equipment for connection to ADP systems, when it is determined that interchange of equipment between different systems is likely.

Implementation. The original version of this standard became effective February 4, 1983. This revision becomes effective December 18, 1990.

Waivers. This standard is non-mandatory. No waivers are required.

FIPS 111, Storage Module Interfaces, is revised as follows:

Applicability. This standard addresses connection of a disk drive to a controller as part of an ADP system. This standard is recommended for use in the acquisition of disk systems that are

connected to small and medium sized computer systems, when it is determined that interchange of equipment between different systems is likely.

Implementation. This standard became effective May 18, 1985. This revision becomes effective December 18, 1990.

Waivers. This standard is non-mandatory. No waivers are required.

FIPS 130, Intelligent Peripheral Interface (IPI), is revised as follows:

Section 8, Applicability. This standard applies to the connection of computers to storage peripheral device controllers. This standard is recommended for use in the acquisition of magnetic disk drives, optical disk drives, and tape drives to be connected to minicomputer systems, when it is determined that interchange of equipment between different systems is likely.

Section 10, Implementation. This standard became effective December 16, 1987. This revision becomes effective December 18, 1990.

Section 11, Waivers. This standard is non-mandatory. No waivers are required.

FIPS 131, Small Computer System Interface (SCSI) is revised as follows:

Section 8, Applicability. This standard addresses the connection of small computers to peripheral devices with integral controllers. This standard is recommended for use in the acquisition of storage peripherals and small computer systems for office or laboratory use, when it is determined that interchange of equipment between different systems is likely.

Section 10, Implementation. This standard became effective December 16, 1987. This revision becomes effective December 18, 1990.

Section 11, Waivers. This standard is non-mandatory. No waivers are required.

Dated: December 12, 1990.

John W. Lyons,

Director.

[FR Doc. 90-29563 Filed 12-17-90; 8:45 am]

BILLING CODE 3510-CN-M